Exhibit C

aruba
a Hewlett Packard Enterprise company

ESP Data Center Design

Validated Solution Guide

Aruba Solution TME

September 08, 2022

Table of Contents

ESP Data Center Design

This guide helps IT professionals understand the following design considerations for a data center environment:

- Hardware selection
- Software selection
- Topology
- High availability
- Scalability
- Application performance
- Security

NOTE:

For the most up-to-date information on ESP Data Center solutions, refer to the Validated Solution Guide Program

Introduction

The Aruba Networks ESP Data Center uses technology and tools to transform the data center into a modern, agile service delivery platform. Organizations of any size, distributed or centralized, can benefit from streamlined performance and improved network cost-effectiveness using the ESP Data Center.

The Aruba AOS-CX operating system simplifies overall operations and maintenance using a common switch operating system across the campus, branch, and data center. The system can be managed on-premises or in the cloud. AOS-CX employs robust artificial intelligence functions that continually analyze and realign network flow to ensure that the system operates seamlessly in accordance with network management best practice, without requiring manual IT staff intervention.

The use of converged Ethernet has changed the way hosts access storage within the modern data center. Dedicated storage area networks (SANs) are no longer required. Lossless Ethernet and bandwidth management protocols ensure timely reads and writes using a traditional IP LAN. The combined cost savings and operational simplicity are driving a major conversion to converged Ethernet.

At the same time, network topologies have become virtualized. Although virtualization delivers the flexibility required to meet the changing data center requirements, it can present complexity and challenges with implementation and management. The Aruba ESP Data Center addresses these challenges by automating installation and implementation of the Aruba AOS-CX operating system, with features such as automated device group configuration, Zero Touch Provisioning, scheduled configuration backups, dashboard-ready network performance metrics, and built-in alerts for critical network functions.

Securing applications and hosts in a data center is critical for maintaining application availability, data integrity, and business continuity. New threats such as ransomware, data exfiltration, and denial of service continue to emerge. Policy and security enforcement requires many tools applied at different layers. The new Aruba CX10000 series switch with Pensando introduces an industry-first distributed services data center switch, capable of performing inline firewall services at wire-speed in the switch itself, focusing on the high level of east-west traffic typical in a data center environment.

Before designing a new or transformed data center, it is important to consider the organization's current and projected strategy for hosting and accessing applications from the cloud. Determine the applications that will remain on-premises so you can establish the data center with ample storage to meet requirements.

To accommodate growth and future adaptation of the network, implementation of a spine-and-leaf underlay that supports software-defined overlay networks is highly recommended. The Aruba Networks CX 83xx, 84xx, and 10000 switching platforms provide a best-in-class suite of products featuring a variety of high-throughput port configurations, industry-leading operating system modularity, real-time analytics, and "always up" performance.

Purpose of This Guide

This guide describes the Aruba ESP Data Center Network deployment, with reference for architectural options and associated hardware and software components. It explains the requirements that shaped the design and the benefits it provides. It introduces Aruba data center solutions that support options for both distributed and centralized workloads. It delivers best practice recommendations to deploy a next generation spine-and-leaf data center fabric using VXLAN and BGP EVPN.

This guide assumes the reader has an equivalent knowledge of an Aruba Certified Switching Associate.

Design Goals

The overall goal is to create a high-reliability, scalable design that is easy to maintain and adapt as business needs change. Solution components are limited to a specific set of products required for optimal operations and maintenance.

Key features of the Aruba ESP data center network include:

- Zero downtime upgrades
- High throughput
- Security
- Converged storage networking
- Flexible segmentation
- Third-party integration

This guide can be used to design new networks or to optimize and upgrade existing networks. Not intended as an exhaustive discussion of all options, the guide focuses on commonly recommended designs, features, and hardware.

Audience

This guide is written for IT professionals who need to design an Aruba ESP data center network. These IT professionals serve in a variety of roles:

- Systems engineers who require a standard set of procedures to implement network solutions
- Project managers who create statements of work for Aruba implementations
- Aruba partners who sell technology or create implementation documentation.

Customer Use Cases

Data center networks change rapidly. The most pressing challenge is to maintain operational stability and visibility for users while moving or upgrading computing and storage resources. In addition, data center teams must continue to support the rapid pace of DevOps environments and meet growing requirements to connect directly and continue operations within the public cloud infrastructure.

Within a rapidly changing landscape, it is critical that network and system engineers responsible for meeting data requirement have efficient tools to streamline and automate complex infrastructure configurations.

This guide discusses the following use cases:

- Pay-as-you-grow designs that support network and computing workload elasticity
- Ease of use and agility to deploy and manage workloads quickly by orchestrating computing, hypervisor, and network management functions
- Improved operations with data center visibility from the computing host to the overall network infrastructure
- Workload mobility, security, and multi-tenancy using standards-based overlay technologies
- Network infrastructure automation and management
- Data aggregation and pre-processing.

Customer Use Cases

Aruba ESP Data Center Network Design

The Aruba Edge Services Platform (ESP) Data Center provides flexible and highly reliable network designs to ensure efficient, reliable access to applications and data for all authorized users while simplifying operations and accelerating service delivery.

Aruba ESP data center includes the following key features and capabilities:

- Modern connectivity — Design efficient and scalable networks to use the full range of port densities and speed options available in the Aruba CX 8xxx and CX 10000 switch families.
- Automation — Automated fabric configuration makes building high-performance, scalable data center networks more efficient and less error-prone.
- Analytics — On-box and cloud analytics ensure that alerts are never missed and that intermittent failures are diagnosed quickly.
- Storage networking — Advanced protocols enable lossless Ethernet with efficient bandwidth reservation and congestion management.
- Host integration — Virtual network visualization is built into the physical network topology for end-to-end management.

The Aruba ESP data center network design may contain one or more of the following elements:

- Aruba Central
- Aruba Fabric Composer
- Aruba NetEdit
- Pensando Policy and Services Manager
- Aruba CX 10000 Ethernet switches with Pensando
- Aruba CX 8xxx Ethernet switches
- Aruba CX 6xxx Ethernet switches for out-of-band (OOB) network management
- Aruba integration into HPE solutions

Aruba ESP Data Center Network Design Options

The Aruba ESP data center supports centralized and distributed workloads anywhere within an organization. Each design supports host uplink bundling, providing throughput and resiliency for mission-critical workloads. Layer 2 domains can be deployed flexibly to meet application requirements and provide virtual host mobility.

Aruba CX switches provide a robust platform for Layer 3 services in the data center. When deployed in a spine-and-leaf topology, a Layer 3 data center network eliminates the need for loop-avoidance protocols, and it is optimized for high capacity and non-oversubscribed, low-latency performance.

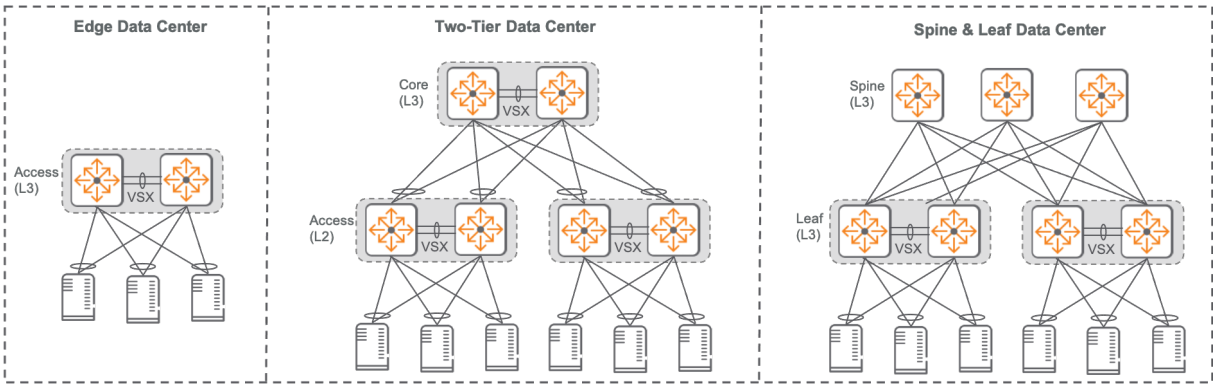

Figure 1: Aruba Data Center Designs

Edge Data Center Overview

Enterprises that have migrated most of their workloads to the cloud and no longer require an on-premises data center can use their existing campus network wiring closets or small server rooms to deploy workloads at the edge.

The same AOS-CX switches that provide wired connectivity to users and Internet of Things (IoT) devices can be leveraged to provide server access.

The edge data center also supports high-bandwidth and low-latency access to computing and storage resources for distributed workloads that may not be well suited to cloud deployments.

Two-Tier Data Center Overview

Enterprises with significant, existing on-premises workloads spanning multiple workgroups often require a traditional, two-tier data center design. The two-tier approach ensures sufficient bandwidth and reliability using traditional protocols such as Link Aggregation Control Protocol (LACP), Spanning Tree Protocol (STP), and Open Shortest Path First (OSPF). Hosts are dual-homed to two top-of-rack (ToR) switches using a Virtual Switch Extension (VSX) multi-chassis link aggregation group (MC-LAG). Each ToR switch is dual-homed to the core. Loops are primarily prevented by using LACP to aggregate redundant links.

Spine-and-Leaf Data Center Overview

Enterprises with growing on-premises workloads and those with workloads spread across data centers can benefit from the efficiency of a Clos-based spine-and-leaf architecture. In most cases, migration to the spine-and-leaf design should be paired with the implementation of a Virtual Extensible LAN (VXLAN) overlay topology. The spine-and-leaf design ensures high reliability using redundant Layer 3 links between leaf nodes and spine switches. Equal-cost multipath (ECMP) routing ensures load balancing and fast failover if a link or switch goes down.

The fully meshed architecture enables simple, horizontal growth by adding another spine switch as needed. VXLAN provides a Layer 2 over Layer 3 tunneling solution, which enables customers to modernize the underlay while preserving legacy service requirements by allowing for physically dispersed Layer 2 segments in the overlay. VXLAN also enables highly segmented designs, which can go beyond traditional VLANs when creating secure, discrete groups of resources within the data center.

This guide addresses the most common use cases of an Aruba spine-and-leaf data center network. For more complex projects not covered in this guide, contact an Aruba or partner SE for design verification.

Aruba ESP Data Center Architecture for Spine and Leaf

Aruba ESP is an evolution of Aruba's end-to-end architecture, providing a Unified Infrastructure with centralized management leveraging Artificial Intelligence Operations (AIOps) for an improved operational experience that helps enable a Zero Trust Security policy. Aruba ESP is the industry's first platform purpose-built for the new requirements of the Intelligent Edge.

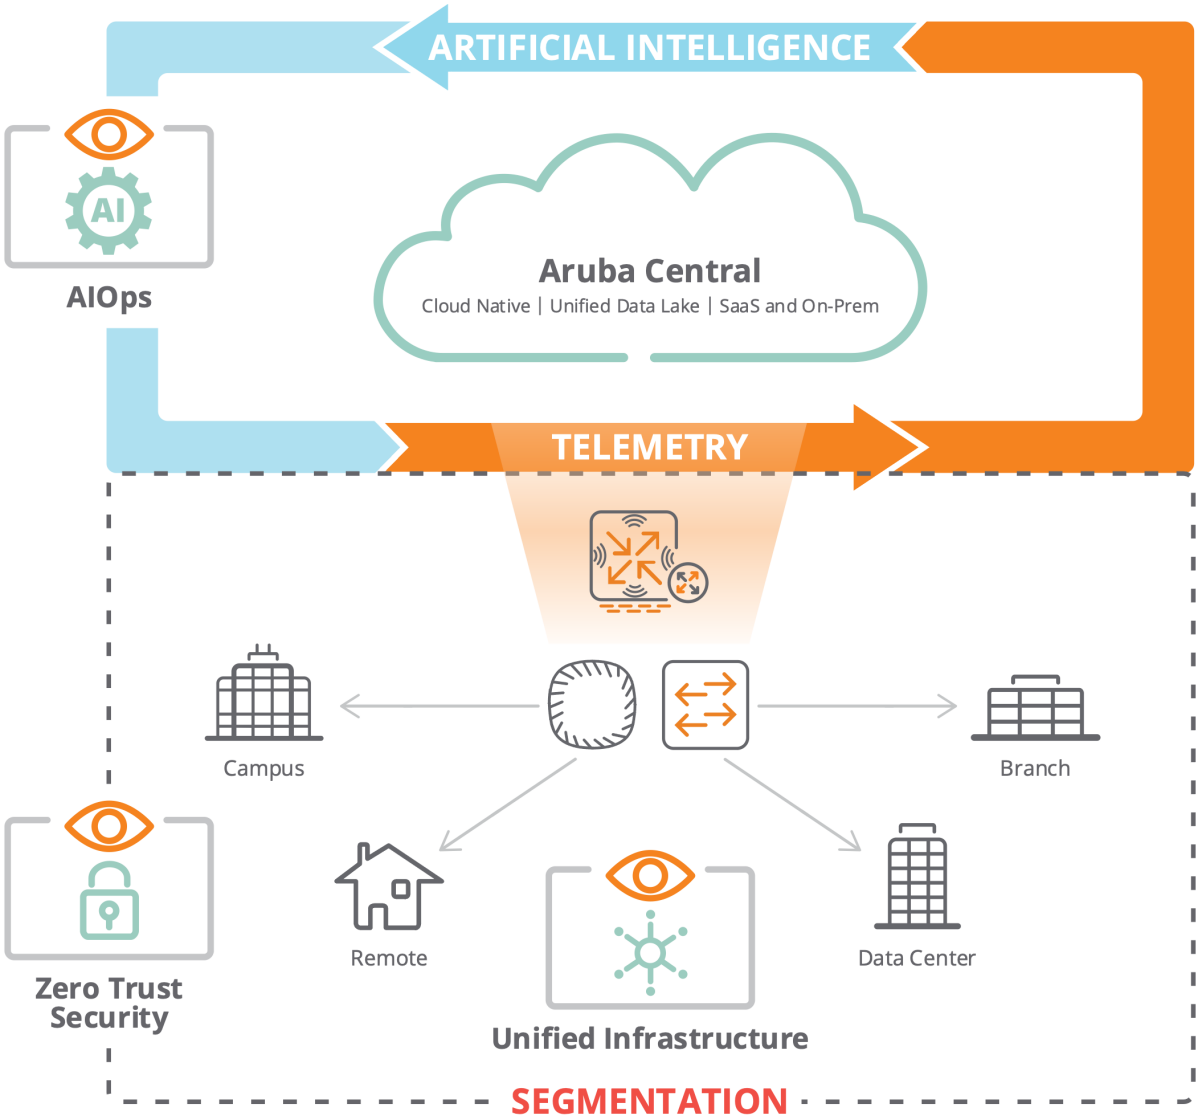

Figure 2: ESP Architecture

Aruba ESP Architecture Layers

Aruba ESP offers a breadth of services, including onboarding, provisioning, orchestration, security, analytics, location tracking, and management. AI Insights reveal issues before they impact users. Intuitive workflow-centric navigation enables the organization to accomplish tasks quickly and easily with views that present multiple dimensions of correlated data. Policies are created centrally, and features such as Dynamic Segmentation allow the network administrator to implement them over an existing infrastructure. The Aruba ESP architecture is built in distinct layers, as shown in the figure below.

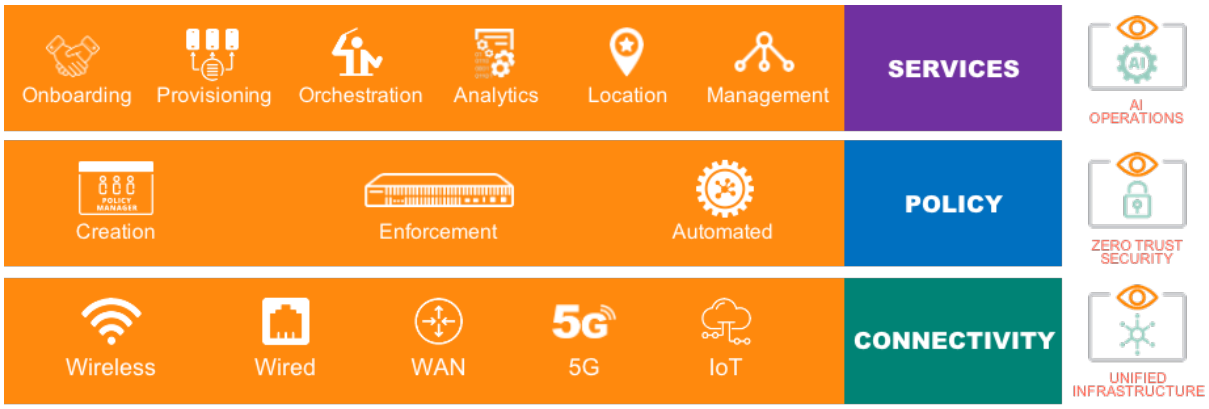

Figure 3: ESP Layers

Aruba ESP Data Center Connectivity Layer

The connectivity layer for the Aruba ESP data center is implemented on Aruba CX 8xxx and 10000 series Ethernet switches, which provide low latency and high bandwidth on a fault-tolerant platform designed to carry data center traffic.

Underlay Network

The underlay network is implemented using a spine-and-leaf fabric topology. It is deployed as a Layer 3 routed network. Each leaf is connected to each spine over a routed port using OSPF as the routing protocol. Layer 2 services are not required in the underlay but can be provided for workloads using virtual overlay networks. The spine-and-leaf underlay topology optimizes performance, increases availability, and reduces latency because each leaf is never more than one hop across multiple load-balanced paths to all other leaf switches.

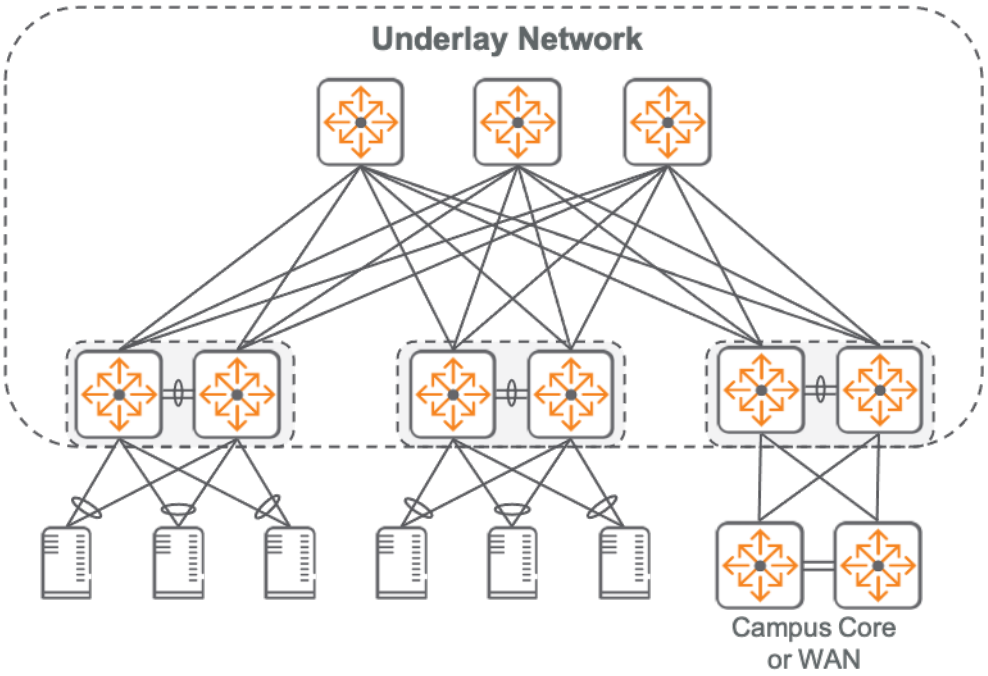

Figure 4: Underlay Network

The spine-and-leaf topology provides a flexible, scalable network design that can expand to accommodate a growing data center without disrupting the existing network. It is easy to begin with a small one- or two-rack fabric that can increase capacity without requiring replacement of existing hardware. ToR ports on leaf switches are used to add computing capacity to a rack incrementally. Ports on spine switches are used to add additional racks to the fabric.

The maximum size of the fabric is determined by the port density on the spine, which is an important consideration for supporting future growth. A minimum of two spine switches is recommended for any size fabric to provide high availability and fault tolerance. Additional spine switches increase overall fabric capacity and reduce the fault domain if a spine must be taken out of service.

Aruba ESP Data Center Policy Layer

The policy layer for the Aruba ESP data center is implemented using overlay technologies and traffic filtering mechanisms to isolate user and application traffic.

Overlay Network

An overlay network is implemented using VXLAN tunnels that provide both Layer 2 and Layer 3 virtualized network services to workloads directly attached to the leaf switches. Similar to a traditional VLAN ID, a VXLAN Network Identifier (VNI) identifies an isolated Layer 2 segment in a VXLAN overlay topology. Symmetric Integrated Routing and Bridging (IRB) enables the overlay networks to support contiguous Layer 2 forwarding and Layer 3 routing across leaf nodes.

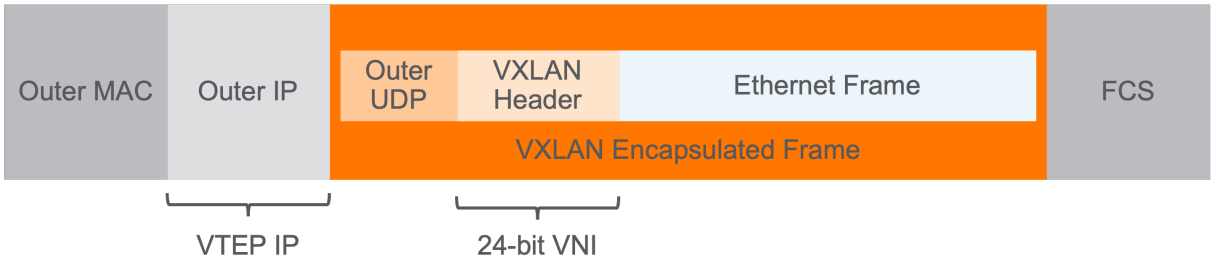

Figure 5: VXLAN Frame

A VXLAN Tunnel End Point (VTEP) is the function within leaf switches that handles the origination and termination of point-to-point tunnels forming an overlay network. A single logical VTEP is implemented when redundant leaf switches are deployed in a rack. Spine switches provide IP transport for the overlay tunnels but do not participate in the encapsulation/decapsulation of VXLAN traffic.

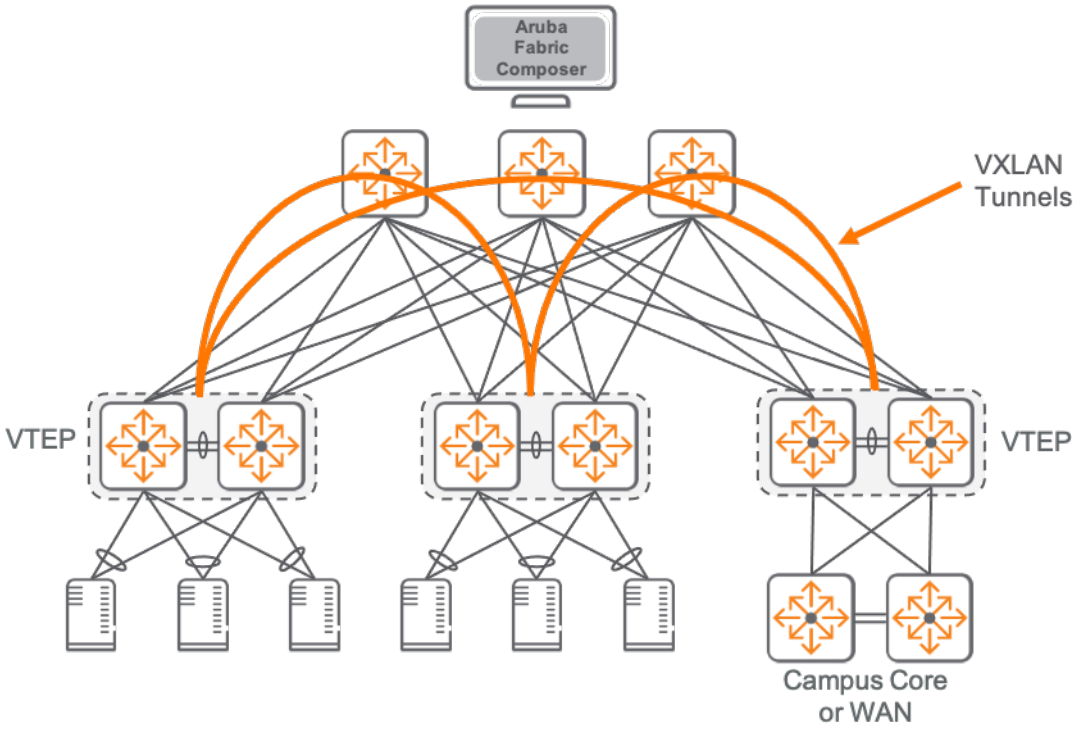

Figure 6: Overlay Network

Attached hosts are learned at the leaf switch using Ethernet link layer protocols. Remote learning across the VXLAN fabric is accomplished using Multiprotocol Border Gateway Protocol (MP-BGP) as the control plane protocol and a dedicated Ethernet virtual private network (EVPN) address family for advertising host IP and MAC prefixes. This approach minimizes flooding while enabling efficient, dynamic discovery of remote hosts within the fabric.

Security and Segmentation

In a VXLAN spine-and-leaf design, a pair of leaf switches is the single entry and exit point to the data center. This *border leaf* is not required to be dedicated to that function. Computing hosts and firewalls also can be attached. Typically, the border leaf is where a set of policies is implemented to control access into the data center network. These policies are the first layer of security for data center applications. They limit access only for permitted networks and hosts while monitoring those connections. The data center perimeter is usually protected in one or both of the following ways:

- Border leaf ACLs — When IP subnets inside the data center are designed to map to security groups or business functions, Access Control Lists (ACL) at the border leaf can provide policy enforcement from user locations into data center applications. If subnets cannot be mapped to security groups, the ACLs can become difficult to manage and scale in larger environments. The primary benefit of perimeter ACLs is that they can be implemented directly on the switching infrastructure to enforce a policy foundation from which to establish data center access. Policies implemented using switch ACLs specifically target Layer 3 and Layer 4 constructs. Switch ACLs are not stateful or application-aware.
- Perimeter firewalls — Dedicated security systems at the perimeter can offer advanced monitoring, application-aware policy enforcement, and threat detection. Perimeter firewalls typically are deployed in transparent or routed mode. In transparent mode, the firewalls behave like a bump in the wire, meaning all allowed user and network control traffic pass transparently through them. In routed mode, a firewall participates in the routing control plane and can be deployed in a configuration that limits the amount of traffic subject to deep inspection. It is important to note that stateful firewalls require symmetric forwarding to apply policy correctly to the subsequent flow.

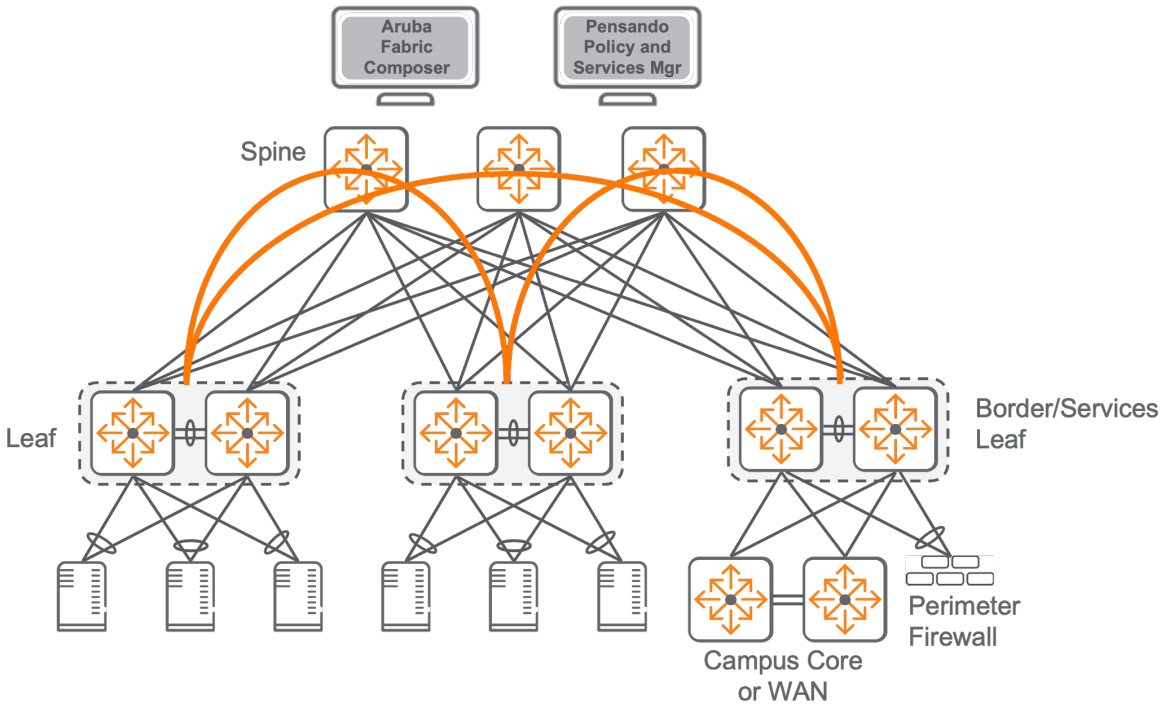

Figure 7: Data Center Policy

Policy inside a VXLAN spine and leaf data center can be implemented using two methods at the network layer: inline using distributed services switches (DSSs) or centrally using a firewall appliance in a services leaf.

- Distributed Services Switch Policy Enforcement — The Pensando programmable data processing unit (DPU) extends Aruba CX 10000 series switches to include stateful firewall capabilities. Using this built-in hardware feature, firewall enforcement is provided inline as part of the switch fabric. There are several advantages to this approach. Firewall policy can be granular to the host with support for microsegmentation. Data center hosts can use local gateways, so east-west traffic flows are optimized between data center hosts. There is no requirement to hairpin data through a services leaf firewall. The Pensando DPU provides wire-rate performance and can alleviate resource consumption on virtualized firewall services processing large data flows by moving firewall services to dedicated switch hardware.
- Services Leaf Policy Enforcement — Another commonly deployed policy enforcement approach is placing a firewall appliance in a services leaf. Firewalls connected at the services leaf are used as the default gateway for hosts requiring specific services accessible through the firewall. An advantage of this approach is the ease with which a Layer 2 overlay network can be used to transport host traffic to the firewall. The disadvantage is that it relies on a centralized gateway and prevents the use of an active gateway at every ToR for optimal forwarding. Similar to a border leaf, the services leaf is not required to be dedicated to this function. The following diagram illustrates the inefficient traffic hairpin, when using a services leaf firewall.

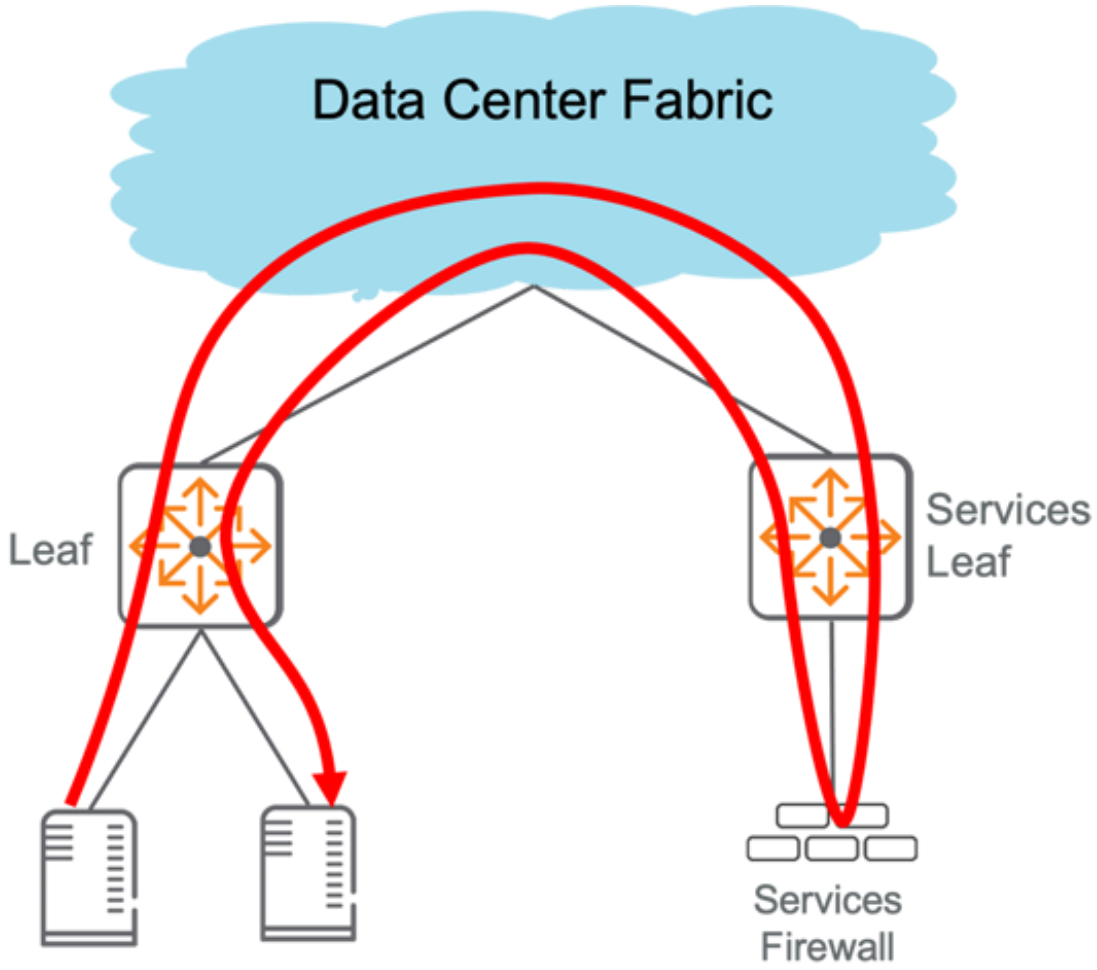

Figure 8: Services Firewall Hairpin Diagram

Some vendors offer virtualized firewall services within a hypervisor environment. This approach can provide granular, service-level policy enforcement while allowing for the use of active gateways. VMware NSX is an example of a product that can integrate in this way. VXLAN overlays can be implemented in both hardware and software to achieve optimal network virtualization and distributed firewall services while securing east–west traffic inside the data center.

Aruba ESP Data Center Services Layer

The Aruba ESP data center solutions include management plane choices that enable an organization to apply the approach that best suits its needs.

- Aruba Central provides a cloud management solution for the end-to-end Aruba ESP solution.
- Aruba Fabric Composer (AFC) is an on-premise fabric automation tool that provides a simplified, workflow-based method of fabric configuration.
- Pensando Policy and Services Manager (PSM) provides management and monitoring of Pensando DPUs contained in Aruba CX 10000 switches.

- Aruba NetEdit provides the same multidevice configuration editor and topology mapper now found in Aruba Central in an on-premises offering.

Aruba Central

Aruba Central is designed to simplify the deployment, management, and optimization of network infrastructure. The use of integrated Artificial Intelligence (AI)-based Machine Learning (ML), and Unified Infrastructure management provides an all-encompassing platform for digital transformation in the enterprise.

Aruba Central provides advanced services to facilitate transformational data center rollouts. NetEdit MultiEditor capability is now integrated into Central, making it possible to deploy complex, multidevice, multilayer configurations from the cloud to the data center. The Network Analytics Engine provides real-time alerts on the state of switches and allows for rapid analysis of intermittent problems. Aruba Central is cloud-hosted for elasticity and resilience, which also means that users need not be concerned with system maintenance or application updates.

Workflow-based configurations within Central allow for efficient, error-free deployments of Aruba solutions anywhere in the world. The workflows are based on common best-practice approaches to network configuration. They enable new devices to be brought online quickly using new or existing network configurations.

AIOps

According to Gartner Glossary, “AIOps combines big data and machine learning to automate IT operations processes, including event correlation, anomaly detection and causality determination.”

Aruba AIOps, driven by Aruba Central, eliminates manual troubleshooting tasks, reduces average resolution time, and automatically discovers network optimizations. Aruba’s next-generation AI uniquely combines network- and user-centric analytics to identify and inform personnel of anomalies. It also applies decades of networking expertise to analyze and provide prescriptive actions.

AI Insights monitor connectivity performance, radio frequency (RF) management, client roaming, airtime utilization, and wired and SD-WAN performance. Each insight is designed to reduce trouble tickets and meet service-level agreements (SLAs) by addressing network connectivity, performance, and availability challenges.

AI Assist uses event-driven automation to trigger the collection of troubleshooting information, identify issues before they impact the business, and virtually eliminate the time-consuming process of log file collection and analysis. After log information is collected automatically, IT staff are alerted with relevant logs that can be viewed and shared with Aruba TAC, who can assist more quickly with root cause determination and remediation.

Aruba Fabric Composer

AFC provides API-driven automation and orchestration capabilities for the Aruba ESP data center. AFC discovers and interrogates data center infrastructure to automate and accelerate spine-and-leaf fabric provisioning as well as perform day-to-day operations across rack-scale computing and storage infrastructure. AFC orchestrates a set of switches as a single entity called a *fabric* and enables the operator to orchestrate data center resources using an application-centric approach to visualizing network and host infrastructure.

Visualization of the data center network fabric includes physical and virtual network topologies as well as host infrastructure through integration with ArubaOS-CX, HPE iLO Amplifier, HPE SimpliVity, VMware vSphere, and other leading data center products. In addition to providing a complete view across the fabric, AFC makes network provisioning accessible to others beside high-level network staff. It provides a secure platform for orchestrated deployment of host and networking resources across the fabric using a guided workflow user interface. AFC ensures a consistent and accurate configuration of a spine-and-leaf data center with or without deployment of an overlay network.

AFC product integration with vSphere facilitates automatic modification of both source and destination IPs in security policy using dynamic endpoint groups. This empowers VMware administrators to add or remove hosts from firewall rules or ACLs by modifying tags associated with a VM guest.

AFC is an end-to-end data center network management tool recommended for new data center deployments based on a spine-and-leaf fabric topology. It is particularly helpful when also deploying an overlay topology using VXLAN-EVPN. AFC configures both the underlay and overlay routing automatically using basic IP information provided by the operator.

AFC facilitates stitching multiple fabrics together to support extending an overlay across multiple locations.

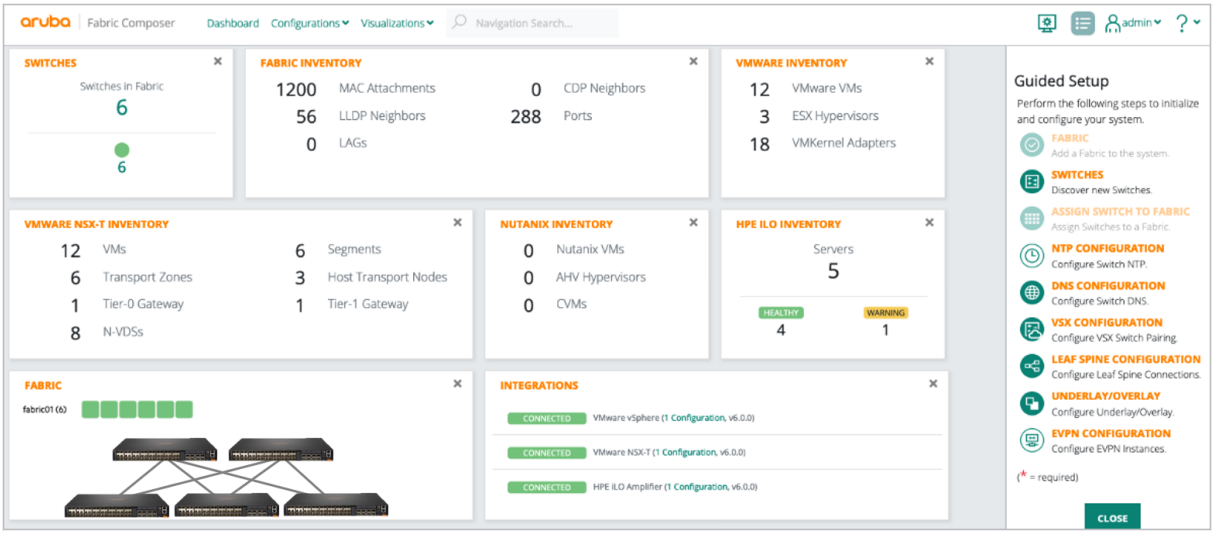

Figure 9: Aruba Fabric Composer

Pensando Policy and Services Manager

The Pensando Policy and Services Manager (PSM) provides an API-based platform for programming and monitoring Pensando DPUs integrated into Aruba CX 10000 switches. PSM is the firewall policy authority for associated switches.

AFC integration with PSM enables single-pane-of-glass configuration and orchestration of both the switch fabric and PSM-managed services.

Aruba NetEdit

Aruba NetEdit helps IT teams automate the configuration of multiple switches and ensure that deployments are consistent, conformant, and error-free. It enables automation workflows without the overhead of programming by providing operators with a user-friendly interface similar to command line. NetEdit also provides a dynamic network topology view to ensure an up-to-date view of the network.

When deploying an Aruba ESP data center network using on-premises tools, NetEdit should be deployed for detailed configuration management. While Aruba Fabric Composer enables fast, error-free spine-and-leaf implementations, NetEdit provides the ability to tailor that configuration when necessary. Together, Fabric Composer and NetEdit deliver an automated, integrated, and validated network configuration ready to support the needs of any data center network.

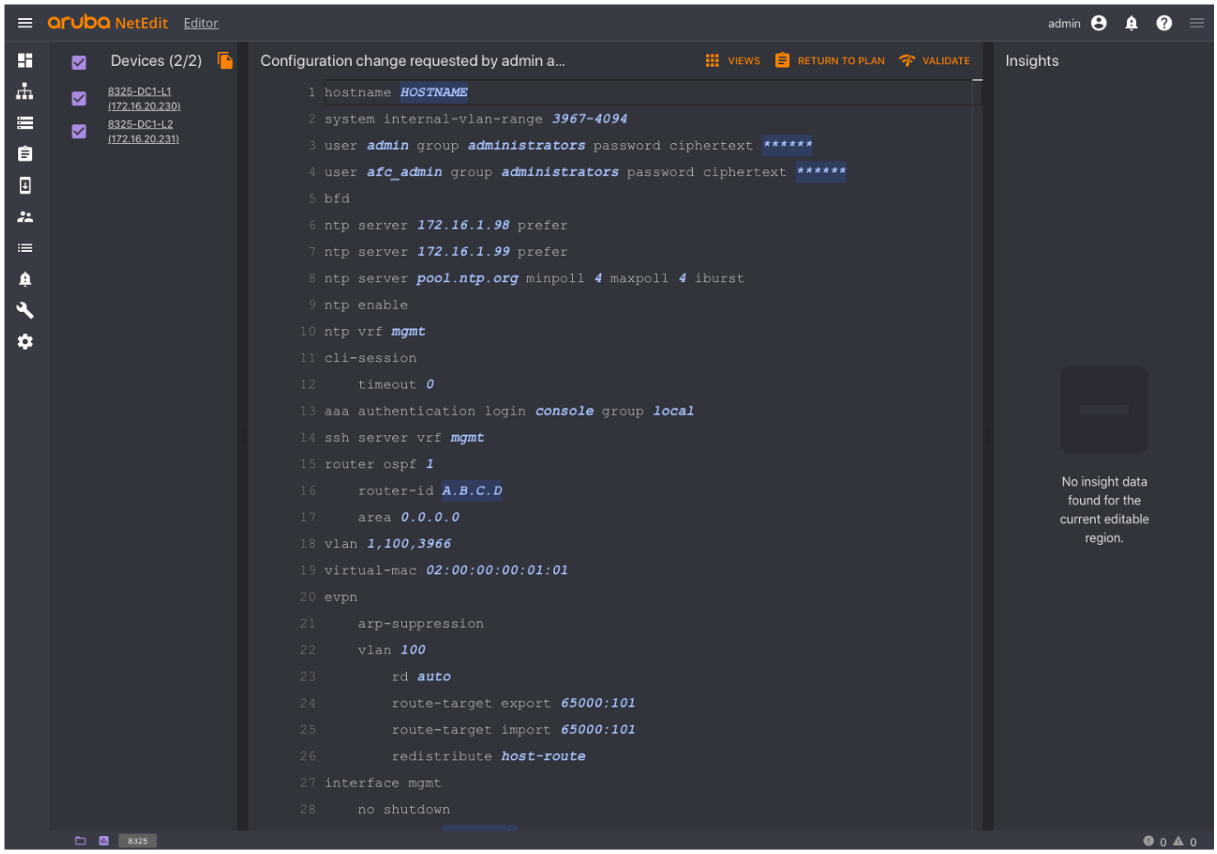

Figure 10: NetEdit

Aruba Network Analytics Engine

Aruba Network Analytics Engine (NAE) provides a built-in framework for monitoring and troubleshooting networks. It automatically interrogates and analyzes network events to provide unprecedented visibility into outages and anomalies. Using these insights, IT personnel can detect problems in real time and analyze trends to predict or even avoid future security and performance issues

A built-in time-series database delivers event and correlation history along with real-time access to network-wide insights to help operators deliver better user experiences. Rules-based real-time monitoring and intelligent notifications automatically correlate to configuration changes. Integrations with Aruba NetEdit and third-party tools such as ServiceNow and Slack provide the ability to generate alerts to trigger actions within the IT service management process.

NAE runs within the AOS-CX operating system in the Aruba CX 6xxx, CX 8xxx, and CX 10000 switch series. NAE agents test for conditions on the switch, its neighboring devices, or on traffic passing through the network, and then take actions based on the result of the test.

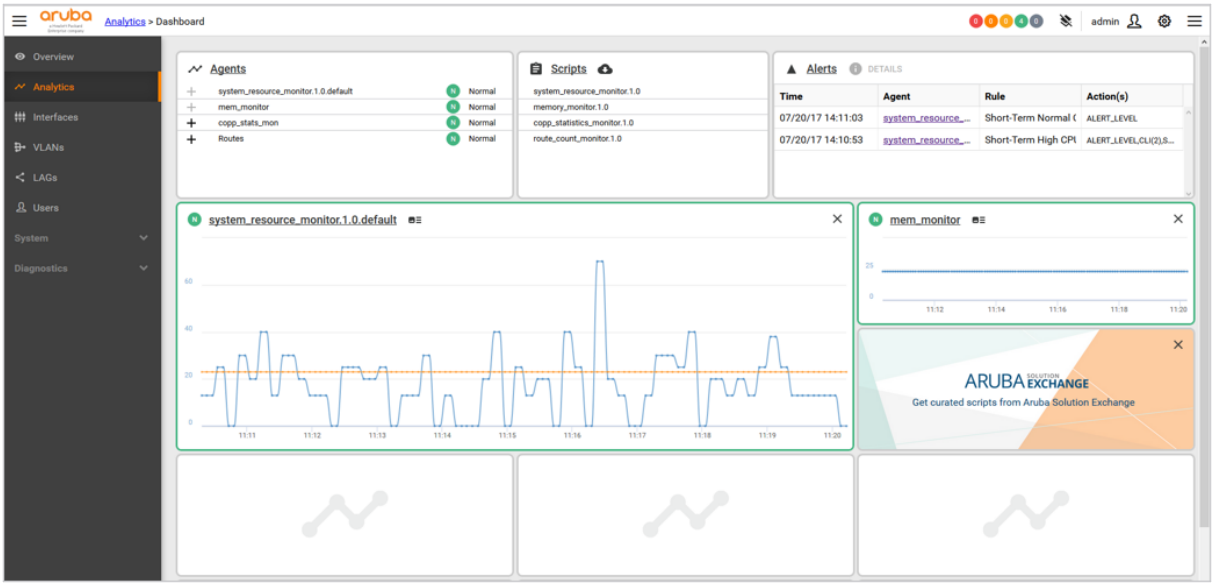

Figure 11: Network Analytics Engine

Choosing an Approach

In general, small, edge-connected data centers are best managed using Aruba Central to ensure consistent configuration anywhere in the world. Larger, centrally located data centers likely require the use of Aruba NetEdit so that detailed, custom configurations can be written and deployed automatically to multiple network devices.

Plans to build a spine-and-leaf data center topology should include AFC. When deploying a VXLAN overlay, AFC is highly recommended to simplify the configuration of underlay and overlay services as well as Layer 3 segments. When deploying PSM with Aruba CX 10000 switches, AFC is recommended to manage firewall rule and policy creation.

Additional Data Center Services

Planning a data center network involves more than just designing the physical network infrastructure. It also is necessary to ensure that services are available to bring switches and hosts online and to ensure that devices can send log messages to a syslog server accessible to people and applications.

It may be useful to leverage the Zero Touch Provisioning (ZTP) capabilities of Aruba switches. To use ZTP, the network must provide a Dynamic Host Configuration Protocol (DHCP) server on a management LAN with a route to the Internet. In addition to the default gateway address, devices also require at least one domain name service (DNS) server to resolve hostnames required for connectivity to Aruba Central and the Aruba Activate service.

Network Time Protocol (NTP) ensures that log data from across the network and in the cloud is time stamped correctly for later analysis. NTP also is required for public key infrastructure (PKI) to function correctly. PKI is required for a variety of access security approaches today. A log management or security information and event management (SIEM) solution also is a part of most modern data centers, which can be used to establish baselines for all switches in the network.

Aruba ESP Data Center Design for Spine and Leaf

To design a successful spine-and-leaf data center, consider all layers of the data center network. This section provides general design considerations for different layers of the Aruba ESP data center. The reference architecture section includes hardware-specific recommendations useful for finalizing the design.

Connectivity Layer Design

Design recommendations and best practice for physical connectivity of computing and network infrastructure are provided below.

Host Connectivity

The first step in designing a data center is identifying the types of connectivity required by the computing hosts. Server hardware typically has an Ethernet RJ45 port for a lights-out management device such as HPE iLO. Applications are commonly connected over redundant links using RJ45 or small form factor pluggable (SFP) ports.

The lights-out port is typically connected using a Cat5e or Cat6 copper patch cable into a switch on the management LAN. The host-to-leaf connections are usually 10 Gb or 25 Gb using SFP+/SFP28 fiber modules, copper direct-attach cables (DAC), or active optical cables (AOC). DACs have limited distance support and can be harder to manage due to the thicker wire gauge when compared with optical cables. AOCs support longer distances than DACs, and AOCs are thinner and easier to manage. Both DACs and AOCs cost less than separate fiber patch cables and optical transceivers, but they operate only at a single speed.

It is important to verify that both the host network interface controller (NIC) and the ToR switch are compatible with the same DAC or AOC. When separate transceivers and optical cables are used, verify transceiver compatibility with the host NIC, ToR switch, and optical cable type. The supported transceiver on the host is often different from the supported transceiver on the switch. Always consult with a structured cabling professional when planning a new or upgraded data center.

When deploying a converged network for IP storage traffic, look for NIC cards that support offload of storage protocols. Similarly, ensure that the hardware also supports VXLAN offloading. These features help minimize latency of storage traffic by reducing the load on a host CPU.

Applications can be hosted directly on a server using a single operating system, commonly referred to as a “bare-metal” server. Multiple hosts can be virtualized on a single physical server using a hypervisor software layer. Examples include VMware ESXi or Microsoft Hyper-V.

Hypervisors contain a virtual switch that provides connectivity to each virtual machine (VM) using Layer 2 VLANs or VXLAN tunnels for segmentation. A successful spine-and-leaf design should support Layer 2 and Layer 3 connectivity using untagged and VLAN-tagged ports to match the required connectivity to the server and/or virtual switch inside the server. AFC provides visibility and orchestration of the configuration between the server and Aruba ToR switches to ensure that connectivity is established properly.

Out-of-Band Management

The Aruba ESP data center spine-and-leaf design uses a dedicated management LAN connecting to switch management ports and to host lights-out management (LOM) ports. Typically, a single management switch is deployed at every rack for OOB management. A dedicated management switch ensures reliable connectivity to data center infrastructure for automation, orchestration, and traditional management access.

Top-of-Rack Design

Deploying switches in the ToR position enables shorter cable runs between hosts and switches. The result is a more modular solution with host-to-switch cabling contained inside a rack enclosure and only switch uplinks exiting the enclosure. This approach helps reduce complexity when adding racks to the data center.

In the Aruba ESP data center, each rack is serviced by a redundant pair of VSX-configured switches. This enables the connection of dual-homed hosts to two physical switches using a link aggregation bundle for fault tolerance and increased bandwidth.

VSX enables a distributed and redundant architecture that is highly available during upgrades. It virtualizes the control plane of two switches to function as one device at Layer 2 and as independent devices at Layer 3. From a data-path perspective, each device does an independent forwarding lookup to decide how to handle traffic. Some of the forwarding databases, such as the MAC and ARP tables, are synchronized between the two devices via the VSX control plane over a dedicated inter-switch link (ISL) trunk. Each switch builds the Layer 3 forwarding databases independently.

When deploying a pair of switches in VSX mode, ensure that three ports connect the switches to one another. Two ports are members of a link-aggregation data path between the switch pair and should be the same speed as the uplinks ports. A third can be any available lower speed port to preserve uplinks for future spine and VSX link-aggregation connectivity.

For backward compatibility and to support future growth, choose a ToR switch that supports connectivity rates of 1, 10, or 25 Gb/s. These connection speeds can be implemented using the same fiber-optic media types so bandwidth can be increased simply by upgrading transceivers or DACs/AOCs.

Keep the following in mind when selecting a ToR switch:

- DSS feature requirements: The Aruba CX 10000 is required in a data center design that uses the inline stateful firewall inspection performed by the Pensando programmable DPU.
- Number of and type of server connections: Typical rack server configurations support 48 host-facing ports, but lower-density ToR options are available in the Aruba CX 8360 series.
- Host connectivity speed: To simplify management, consolidate hosts connecting at the same speeds to the same racks and switches. Adapting the port speed settings of a particular interface between 25 and 10 Gb may impact a group of adjacent interfaces. Consider interface group size when planning for a rack requiring multiple connection speeds.
- Number of uplink ports: ToR switch models support a range of uplink port densities. When using VSX for redundancy, two uplink ports are used for ISLs providing data-path redundancy and cannot be used for spine connectivity.
- ToR-to-spine connectivity: The amount and port speed of the uplinks define the oversubscription rate from the hosts to the data center fabric. For example, in a four-spine deployment at 100 Gb, a non-oversubscribed fabric can be implemented for racks of 40 servers connected at 10 Gb.
- Cooling design: Different ToR models are available for port-to-power and power-to-port cooling. In power-to-port configurations, an optional air-duct kit can isolate hot air from servers inside the rack. Cabling can absorb heat and restrict airflow. Short cable routes and good cable management improve the airflow efficiency.

Spine Design

The spine layer provides aggregation for leaf switches. In a spine-and-leaf design, each ToR switch is connected to each spine switch. Each leaf-to-spine connection should use the same link speed to ensure multiple equal-cost paths within the fabric. This enables ECMP-based routing to ensure connectivity if a link goes down.

The port capacity of the spine switches defines the maximum number of racks the data center can connect. For a redundant ToR design, the maximum number of racks is half the port count on the spine switch. Two spines is the recommended minimum for high availability. Additional spines increase overall fabric capacity and reduce the size of the fault domain if a spine is out of service.

- Determine rack media and bandwidth requirements.
- Determine if single or redundant ToR switches will be installed.
- Determine how many racks are needed for current computing and storage requirements.
- Determine the spine switches required to support the planned racks.

- Design the data center network for no more than 50% capacity to leave room for growth.

If the network has more than two spine switches, pay attention to the number of uplink ports available on the chosen ToR switch. Each ToR switch must connect to each spine for ECMP to work effectively.

When deciding where to place the spine switches, consider their distance from the leaf switches and the media type used to connect them. Leaf-to-spine connections are either 40 Gb or 100 Gb fiber using quad SFP (QSFP) transceivers or AOCs, in which the cable and transceiver are integrated, similar to DACs.

Underlay IP Design

The underlay of a spine-and-leaf data center network is the layer that provides IP connectivity between spine-and-leaf switches. The network underlay ensures that VXLAN tunneled traffic (the overlay network) can be forwarded across the fabric.

The Aruba ESP data center uses OSPF as the underlay routing protocol. OSPF is a widely used, well understood Interior Gateway Protocol (IGP) that provides straightforward configuration and fast convergence. A single OSPF area and point-to-point interfaces are recommended to minimize the complexity and time required to establish neighbor adjacencies.

Configure the data center switches for a jumbo maximum transmission unit (MTU) of 9198 bytes. This accommodates the storage protocols likely to be deployed and the expanded frame header used by VXLAN.

Policy Layer Design

This section provides design recommendations and best practices for the policy layer design of the data center network.

Out-of-Band Management Network

Organizations should plan to build a physically separate management LAN and role-based access control on the network devices. This means that login to a switch requires authentication against an enterprise directory, typically accomplished using the TACACS+ protocol and a policy server such as Aruba ClearPass Policy Manager. Use of logging facilities, log management, and log analysis also should be considered.

Establish a separate management network to ensure that data center switch reachability is not blocked unintentionally when modifying data plane policy.

Segmentation and Policy Summary

Segmentation is a logical separation of data center hosts between which a security policy can be enforced. Network segments can be created between hosts internally in the data center and between the data center and external networks.

Security policy specifies the type of traffic allowed between network segments. Network-based policy is typically enforced using a firewall or ACL. If traffic is permitted by a stateful firewall, dynamic state is created to permit return traffic for the session created by the initiating host. An ACL is applied to traffic in only one direction with no dynamic state creation.

Applying network security policy plays a significant role in reducing the attack surface exposed by data center hosts, constraining the options for lateral threat movement if a host has been compromised and preventing data exfiltration.

Blocking unnecessary protocols reduces the available tactics a threat actor can use in host exploitation. Blocking can be applied to both north-south and east-west data center traffic.

Scoping allowed outbound traffic inhibits command and control structures and blocks common methods of data exfiltration. As east-west data center traffic has grown to comprise the majority of traffic for data center hosts, applying intra-data center security policy is critical for maximizing security.

Segmentation of data center hosts into distinct routing domains is a key method of scoping internal data center communication. Network hosts that are members of one routing domain are not capable of communicating with hosts in another routing domain by default. Each routing domain contains a set of hosts that must communicate with each other, while traffic into or out of the routing domain can be controlled. Multiple strategies are available to share connectivity between routing domains, when necessary.

A single switch can contain multiple routing domains when supporting virtual routing and forwarding instances (VRFs). Each VRF is its own independent routing domain. A VRF instance can correlate with a customer, an application, a set of hosts with common security requirements (i.e., PCI), or a segment of the network (i.e., production services environment, development environment, etc.). Each VRF consists of a unique route table, member interfaces that forward traffic based on the route table, and routing protocols that build the route table. A VRF may contain overlapping IP addresses with another VRF, because the individual route tables are discrete.

VRFs should be added in a thoughtful manner, as the operational complexity of a network increases as the number of VRFs grows. Minimizing complexity results in a quickly identifiable network structure that is easier to troubleshoot. Each organization should define its own policy that clearly states the criteria to be met when adding a VRF to the network. For example, a service provider that supports multiple tenants may require many more VRFs than a university data center.

VRF member interfaces that connect to external networks provide a natural point for implementing security policy.

Segmentation also refers to the separation of hosts by VLAN or private VLAN. Traffic between VLANs must be routed, and all host traffic between VLANs is forwarded via an IP gateway interface, where security policy can be applied. A private VLAN allows for more granular control of communication between individual hosts even within the same VLAN on the same switch. The Aruba ESP data center uses private VLANs coupled with the CX 10000 switch as part of a microsegmentation strategy to enforce security policy between individual hosts at the network level, including between VM guests installed on the same hypervisor.

Overlay Control Plane Design

Host mobility refers to the ability to move physical or virtual hosts in a data center network without changing the host network configuration. Especially powerful for virtualized hosts, this flexibility ensures optimized computing resources, high availability of applications, and efficient connectivity for distributed workloads.

To maintain a data center overlay and successfully forward traffic through it, VTEPs within the fabric require reachability information about fabric connected endpoints. A distributed, dynamic control plane is recommended for the following reasons:

- Traditional flood-and-learn techniques can consume large amounts of bandwidth due to the replication of traffic in a large spine-and-leaf environment.
- Network configuration is simplified as the ToR switches automatically learn about other ToR switches inside the fabric.
- A distributed control plane provides redundancy and a consistent topology state across the data center fabric switches.
- A distributed control plane allows optimal forwarding using distributed gateways at the ToR switches. This makes it possible for the default gateway address to remain the same across the fabric.

The use of MP-BGP with EVPN address families between VTEPs provides a standards-based, highly scalable control plane for sharing endpoint reachability information with native support for multi-tenancy. For many years, service providers have used MP-BGP to offer secure Layer 2 and Layer 3 VPN services on a very large scale. Network operations are simplified by using an iBGP design with route reflectors so that peering is required only between leaf switches and the spine. Some of the more notable BGP control plane to become familiar include:

- **Address Family (AF)** — MP-BGP supports exchanging network reachability information for multiple address types by categorizing them into address families (IPv4, IPv6, L3VPN, etc.). The Layer 2 VPN address family (AFI=25) and the EVPN subsequent address family (SAFI=70) are used to advertise IP and MAC address information between MP-BGP speakers. The EVPN address family contains reachability information for establishing VXLAN tunnels between VTEPs.
- **Route Distinguisher (RD)** — A route distinguisher enables MP-BGP to carry overlapping Layer 3 and Layer 2 addresses within the same address family by prepending a unique value to the original address. The RD is only a number with no inherently meaningful properties. It does not associate an address with a route or bridge table. The RD value allows support for multi-tenancy by ensuring that a route announced for the same address range via two different VRFs can be advertised in the same MP-BGP address family.
- **Route Target (RT)** — Route targets are MP-BGP extended communities used to associate an address with a route or bridge table. In an EVPN-VXLAN network, importing and exporting a common VRF route target into the MP-BGP EVPN address family establishes Layer 3 reachability for a set of VRFs defined across a number of VTEPs. Layer 2 reachability is shared across a distributed set of L2 VNIs by importing and exporting a common route target in the L2 VNI definition. Additionally, Layer 3 routes can be leaked between VRFs using the IPv4 address family by exporting route targets from one VRF that are then imported by other VRFs.

- **Route Reflector (RR)** — To optimize the process of sharing reachability information between VTEPs, using route reflectors on the spines allows for simplified iBGP peering. This design enables all VTEPs to have the same iBGP peering configuration and eliminates the need for a full mesh of iBGP neighbors.

Segmentation Policy Prerequisites

Data center applications are deployed in many different ways. Applications can be implemented as VMs using hypervisors or hosted in bare-metal servers. Containerized apps are highly distributed and usually require connectivity between multiple computing and service nodes. In some cases, a single data center hosts applications for multiple tenants while offering a set of shared services across them. Because applications are deployed with the majority of traffic contained within the data center, it is incorrect to assume that all security threats are external.

Successful data center policy design begins with understanding the requirements of the applications that run in the environment. It is often necessary to re-profile legacy applications when there is not sufficient documentation of the requirements. From a networking perspective, application profiling should document all network connections required for that application to run successfully. These might include connections to backend databases or cloud-hosted services. To properly define policy regarding which connections must be permitted and which must be denied, it is necessary to know the application profile.

Similarly, analyzing the profile of the users accessing the applications and data is typically required. Never leave a data center wide open to a campus, even if it is assumed to be a secure environment. To restrict access, understand the various user profiles associated with the applications and data required. It is important to identify on-campus, remote branch, mobile field workers, and public Internet requirements so that appropriate data center access profiles can be developed to represent their unique requirements.

Virtual Routing Domain Segmentation

Common best practice is to use the minimum number of VRFs required to achieve clearly defined organizational goals, since each additional VRF increases the complexity of a network. VRFs are employed to support the following use cases:

- Separate production and development application environments. This provides a development sandbox while minimizing risk to production application uptime, and it supports overlapping IP space when required.
- Apply policy to segmented traffic that requires strict regulatory compliance, such as PCI or HIPAA.
- Apply policy to segmented traffic from hosts identified by organizational policy as requiring segmentation and possessing a common set of security requirements. These sets of hosts often share a common administrative domain.
- Isolate Layer 3 route reachability in a multi-tenancy data center, while supporting overlapping IP space.

Data center routes likely require sharing with campus network segments. A direct VRF-to-VRF peering between a data center border VRF and its campus VRF neighbor enables IP segmentation to be extended into the campus. This is referred to as VRF-lite. Multiple data center VRFs can peer with the default campus routing instance, when policy enforcement between data center segments is the only requirement. In both cases, policy enforcement is typically achieved by placing a firewall between data center border leaf switches and campus switches.

Inter-VRF Route Forwarding (IVRF) can be used within a data center to share IP prefixes between VRFs. For example, to provide shared services in a data center, a services VRF can be created to offer a common set of resources and IVRF allows reachability between the application and services VRFs.

VLAN Segmentation

In addition to limiting broadcast domain size, a VLAN and its associated IP subnet can be used to group sets of data center hosts by role, application, and administrative domain.

VLANs that are members of the same VRF have Layer 3 reachability between subnets. These Layer 3 boundaries then become a key point of policy enforcement. VLAN ACLs are typically used to enforce a base policy between subnets in a VRF. When more sophisticated policy requirements arise, the common solution is to deploy a centralized firewall and make it the default gateway for application subnets. This generally results in sub-optimal, inefficient traffic patterns. The Aruba ESP data center provides the option of deploying CX 10000 ToR switches which offer hardware-based, Layer 4 firewall capabilities at the host's uplink switch, thereby reducing the need to hairpin traffic.

Microsegmentation

Microsegmentation extends policy enforcement down to the individual workload and network host level. The Aruba CX 10000 series switch provides a complete and consistent microsegmentation strategy that can be applied across a broad set of data center hosts. Similar to a hypervisor-based firewall, the CX 10000 provides the ability to segment between VM guests on the same VM host using a private VLAN (PVLAN) mechanism. The CX 10000 provides a single, data center microsegmentation strategy that supports all hypervisor types (VMware, Microsoft Hyper-V, KVM, etc.) and bare metal servers. Using the CX 10000 in place of a hypervisor-based implementation offloads policy enforcement cycles from a VM host CPU to dedicated switch hardware.

Microsegmentation can be applied to a subset of hosts requiring a high level of scrutiny, or it can be applied more broadly to maximize a data center's security posture.

Using Aruba CX 10000 Policy

The Aruba CX 10000 with Pensando provides a powerful policy enforcement engine. This section provides background information and details on how to implement CX 10000 firewall policy.

CX 10000 Environments

Applying a consistent PSM-based policy across a data center fabric is more easily achieved when all leaf switches are CX 10000. A mixed environment of DSS and non-DSS capable switches is supported; however, administrators must be mindful of these considerations:

- VM and bare metal hosts requiring firewall policy must be cabled to CX 10000 switches.
- Procedures must be created to prevent automated and manual VM guest migration from a CX 10000-connected VM host to a VM host connected to a non-DSS switch, when the VM guest requires DSS-based policy enforcement.
- A combined set of both egress and ingress policies often must be created to achieve defined security goals.

Host Mobility Implications

Ubiquitous host mobility within a fabric requires that all leaf switches support the same capabilities. The stateful firewall security policies supported on a DSS switch are not available on non-DSS switches. VM mobility must be constrained accordingly. For example, when using dynamic tools such as VMware's Distributed Resource Scheduler (DRS), care must be taken to ensure that virtual switch and port group resources are defined to prevent automated movement of a VM guest requiring firewall services to a VM host that is not connected to a DSS switch.

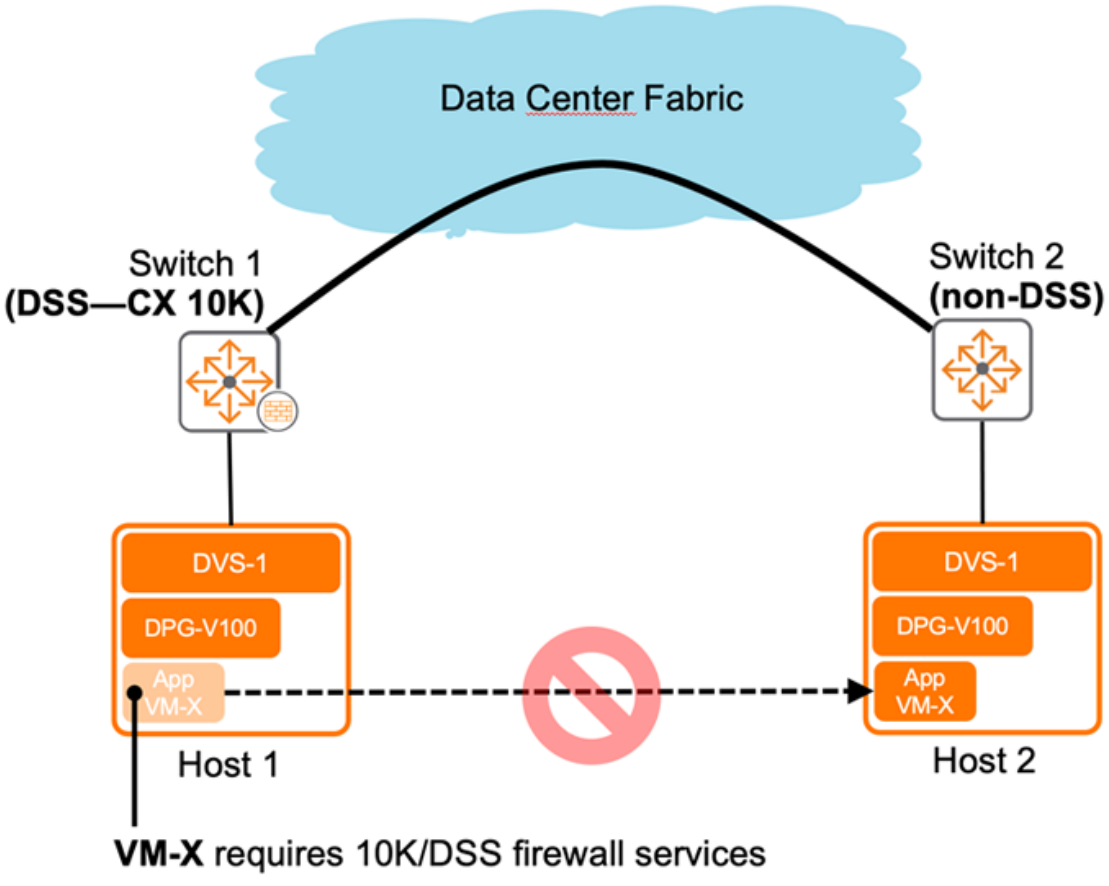

Figure 12: DSS Host Mobility Constraint Diagram

PSM Networks and Policy Basics

PSM associates a *Network* object with a VLAN configured on a CX 10000 switch. Defining a PSM *Network* informs the switch to redirect traffic for the associated VLAN to the Pensando DPU-based firewall for Layer 4 policy enforcement.

PSM assigns policy at two different levels. The *Network* level corresponds with a VLAN. The *Virtual Private Cloud (VPC)* level corresponds to a VRF in an on-premise data center. Policy defined at the VPC level is applied only to VLAN traffic that also has a corresponding *Network* defined.

When traffic for a VLAN (*Network*) is redirected to the Pensando DPU for firewall enforcement, both *VPC* and *Network* policies are applied. PSM evaluates both levels using a logical AND function. If both policies permit the traffic, the traffic is forwarded. If traffic is denied at either level, the traffic is dropped. If a policy is not assigned at one level, policy is enforced by the other level's policy. If no policy is assigned at either level, traffic is permitted. The following matrix summarizes when traffic is allowed or denied.

Network Policy	VPC Policy	Final Result
Permit	Permit	Permit
Deny	Permit	Deny
Permit	Deny	Deny
Deny	Deny	Deny
Permit	No Policy	Permit
Deny	No Policy	Deny
No Policy	Permit	Permit
No Policy	Deny	Deny
No Policy	No Policy	Permit

Figure 13: PSM Policy Decision Matrix

PSM firewall policy is a set of rules that specifies source and destination addresses, and the type of traffic allowed between them using IP protocol and port number. Policy is applied to the *Network* or *VPC* in either an egress or ingress direction, from the perspective of the connected host. Traffic sourced from the host is considered egress, and traffic destined to the host is considered ingress.

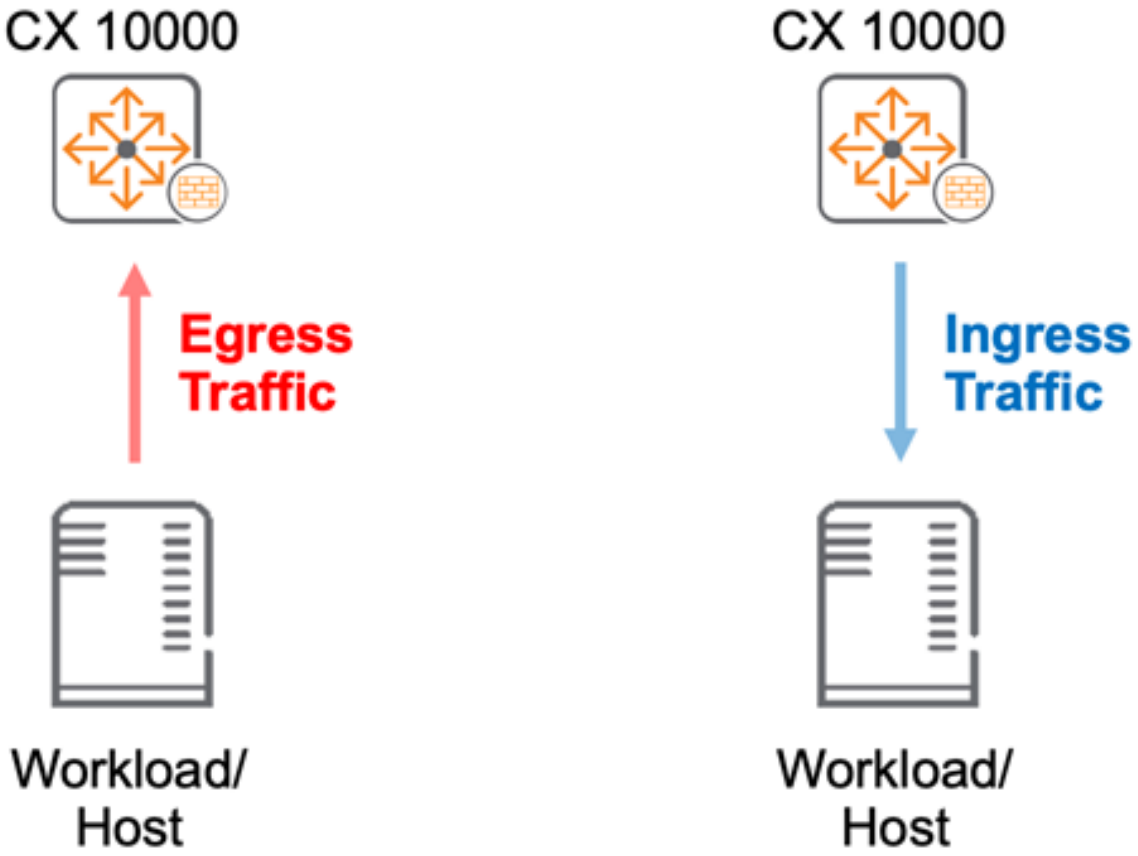

Figure 14: PSM Policy Direction Diagram

An ingress policy applies to traffic from other switches destined to a host in a VLAN with an associated PSM *Network*, where the traffic also hits a logical interface on the destination CX 10000. The logical interface type can be a VLAN SVI or a VTEP. When using an EVPN fabric overlay, ingress policy applies to traffic sourced by hosts that are Layer 2 adjacent in the fabric overlay on other switches. An ingress policy cannot be applied between two hosts attached to the same Aruba CX 10000 switch, even if the traffic hits a VLAN SVI. Ingress policy is not applied to Layer 2 traffic bridged over a traditional Layer 2 link from an adjacent switch.

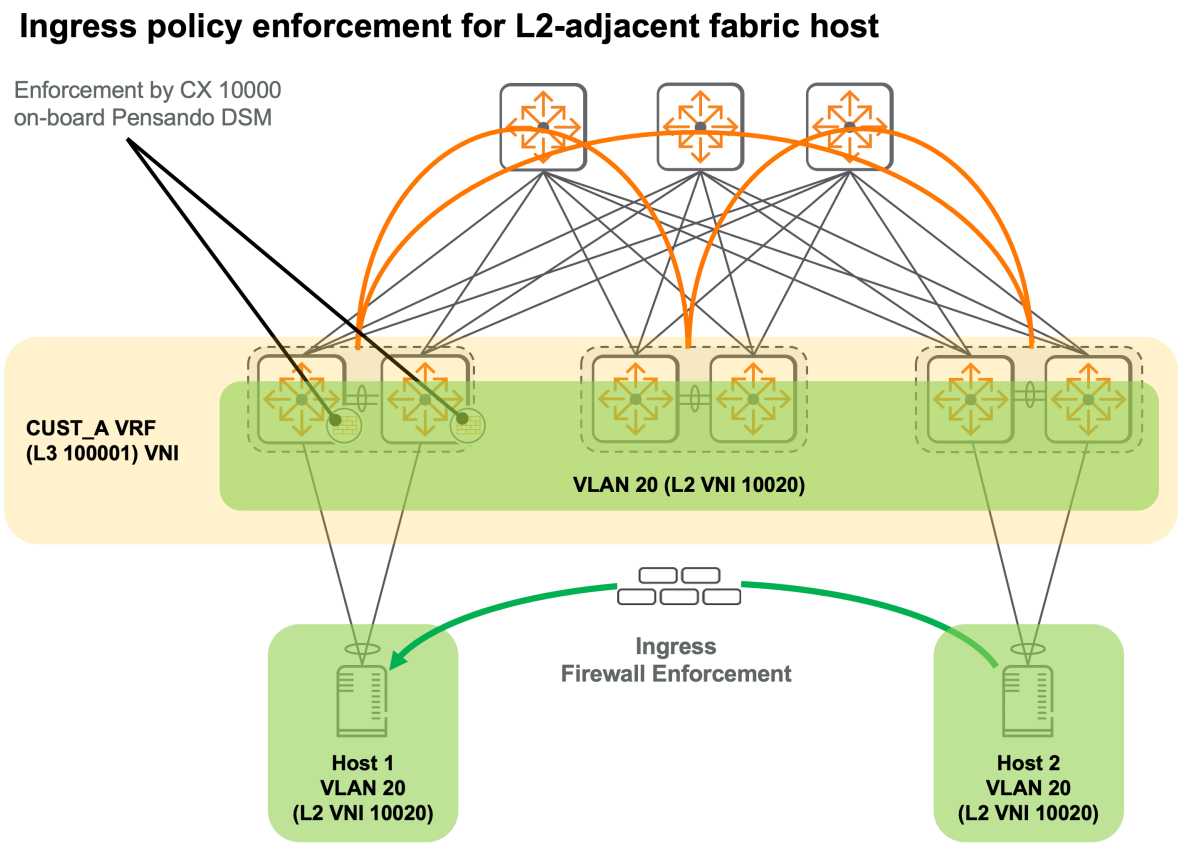

Figure 15: Ingress Fabric Policy Enforcement Diagram

Egress traffic that hits a logical interface is redirected to the Pensando DPU for policy enforcement. In addition to filtering traffic between hosts in the broader fabric context, egress policy also can be applied between hosts attached to the same CX 10000 switch . When applying policy to hosts attached to the same switch, the hosts must reside in separate Networks (VLANs) or be members of a microsegmentation construct. This is required for the traffic to hit a logical interface that can redirect the traffic to the DPU. Traditional Layer 2 bridged traffic is not inspected by egress policy consistent with the same ingress policy restriction.

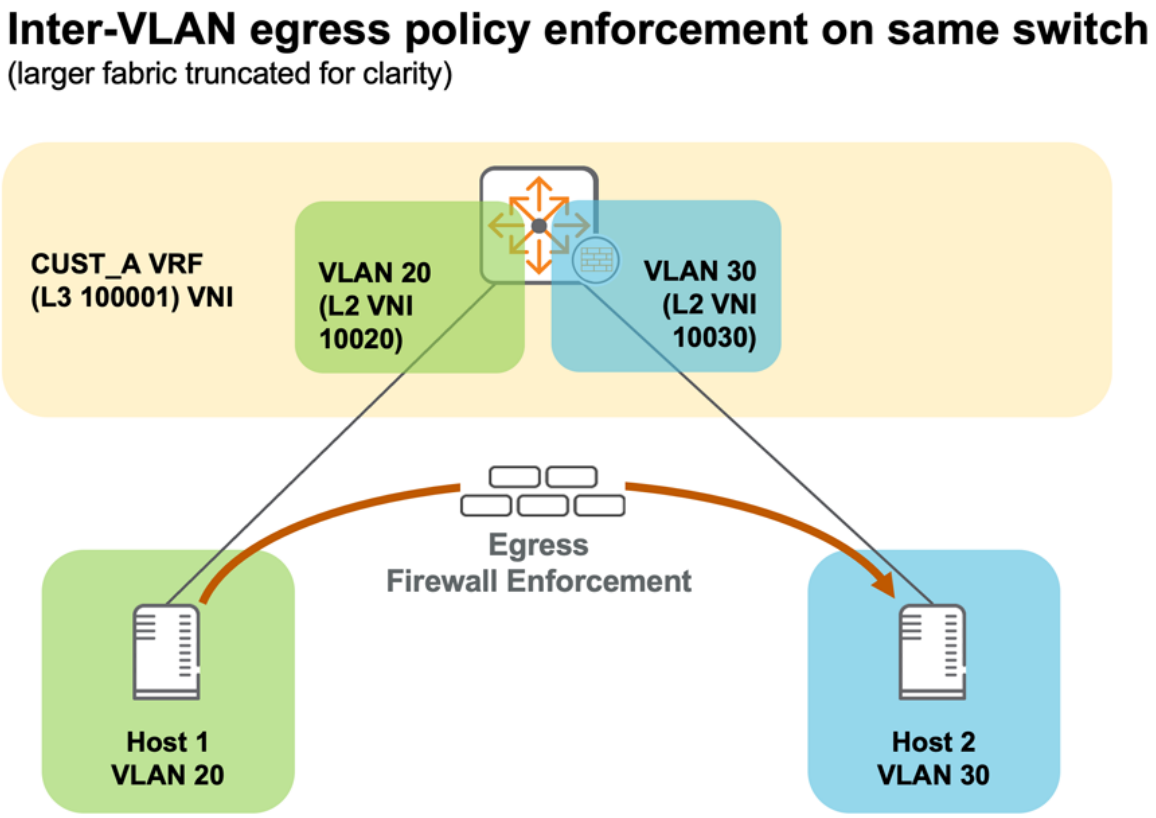

Figure 16: Intra-Switch Egress Policy Enforcement Diagram

Microsegmentation enables the enforcement of PSM firewall policy between hosts in the same VLAN. To perform microsegmentation, a private VLAN (PVLAN) strategy is used to force traffic from attached hosts to hit a VLAN SVI, which allows the switch to redirect traffic to the Pensando DPU for inspection. This is achieved by assigning attached hosts to an isolated PVLAN, so traffic to and from the hosts is allowed only via the isolated PVLAN's associated primary PVLAN. PSM policy is assigned to the PSM *Network* associated with the primary PVLAN. The primary PVLAN SVI uses proxy ARP to respond to ARP requests on behalf of all isolated PVLAN attached hosts. As a result, isolated PVLAN hosts have reachability to each other via the primary PVLAN SVI.

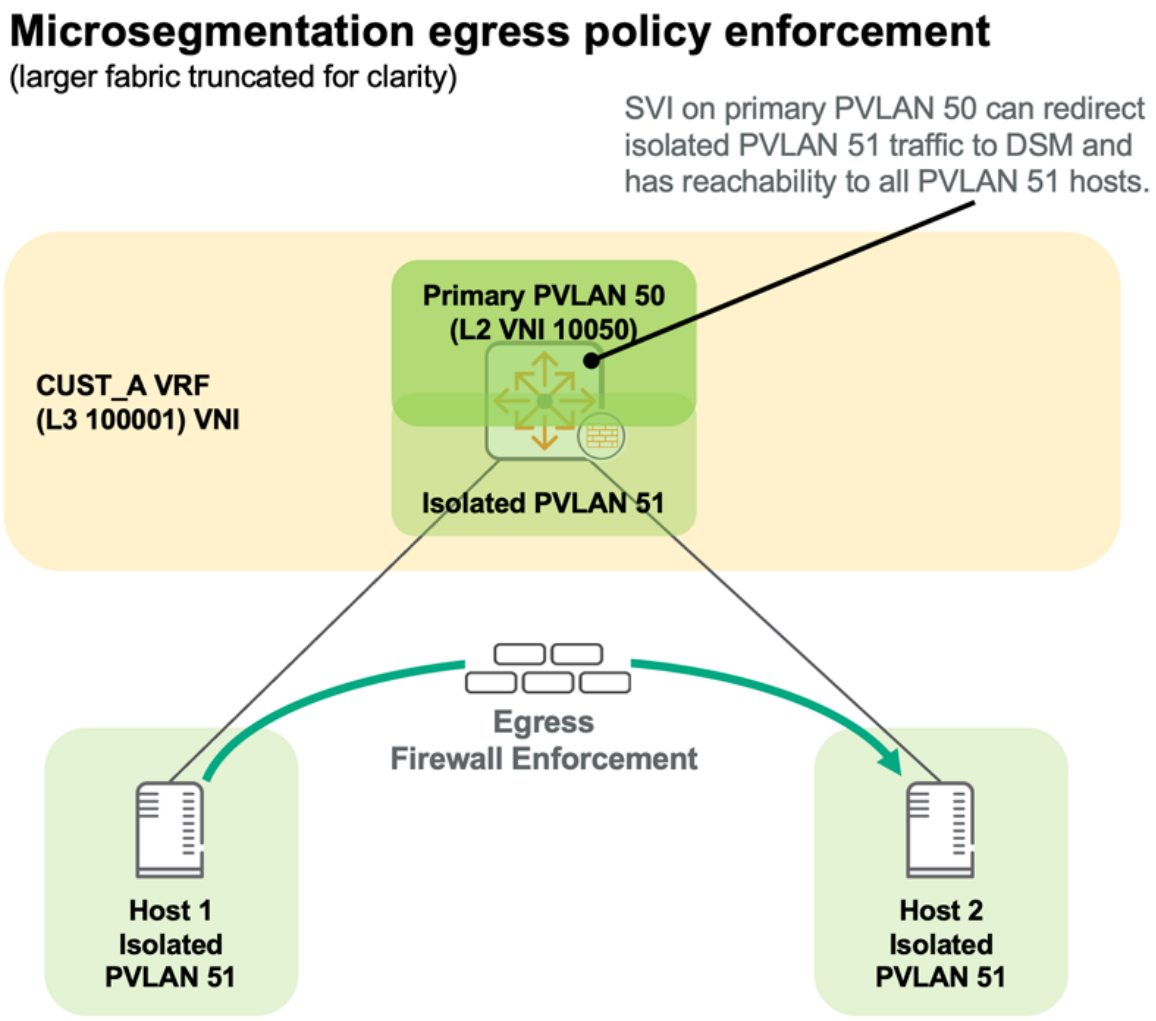

Figure 17: Microsegmentation Egress Policy Enforcement Diagram

In a CX 10000-only environment, egress policy can filter all east-west data center traffic, while leaving the primary role of north-south inbound policy enforcement at the border leaf. In a mixed environment of DSS and non-DSS switches, ingress rules are required for complete east-west data center policy enforcement, which brings inbound sessions from non-data center hosts into scope when defining PSM policy.

It is recommended to apply only one PSM policy level (*Network* or *VPC*) for a particular enforcement direction. The primary advantage of defining policy at the *VPC* level is having a single policy represent the complete rule set for the entire fabric in an enforcement direction. Defining policy at the *Network* level reduces the size of an individual policy and ensures that policy applied to one *Network* does not impact another *Network*.

Mixing policy levels for a single direction can add complexity and duplication of rules in both sets of policy, although mixing policy levels is fully supported. When defining policies at both levels, one policy requires a rule to permit any traffic at the end of the policy, which is recommended in the *Network* level policy. The policy rules above the “permit any” rule should be deny rules. This allows the *VPC* level policy to define what is permitted globally within the fabric with more granular deny restrictions applied at the *Network* level.

PSM Policy Considerations

Defining a PSM *Network* redirects all traffic in and out of the associated VLAN to the Pensando DPU firewall, so network requirements of all VLAN members must be considered in policy rule sets. When an ingress policy is applied, all traffic destined to hosts within the VLAN from other switches must be considered, including all Layer 2 EVPN forwarded traffic and traffic sourced from outside the data center.

When defining an egress policy, all communication sourced by hosts on the VLAN must be considered. Rules allowing underlying services are required when applying an egress policy to a network, because traffic that is not explicitly allowed to be sourced by the hosts is blocked by the implicit deny rule. Rules supporting services such as DNS, logging, and authentication must be defined.

Rules in a policy are applied in the order they appear in the list. An implicit deny all rule is applied at the end of a rule set. Rules that are used more often should appear higher in the list.

It is best practice to define a complete set of rules before applying a policy to a network. If the complete rule set is unknown, an allow-all rule can be applied to collect log data on observed traffic. A complete rule set can be built by inserting rules to allow more specific traffic above the allow-all rule. When no wanted traffic is hitting the allow-all rule at the bottom of the rule set, remove it.

Aruba Reference Architecture for Data Center

The Aruba ESP (Edge Services Platform) data center reference architecture supports high-availability computing racks using redundant top-of-rack (ToR) switches connected in a layer 3 spine-and-leaf topology. The spine-and-leaf topology optimizes performance and provides a horizontally scalable design that can expand to accommodate a growing data center without disrupting existing network components. A data center can start with as few as two spine switches. When additional capacity is required, up to six spine switches can be deployed in a single fabric. The figure below shows the reference architecture with three spine switches.

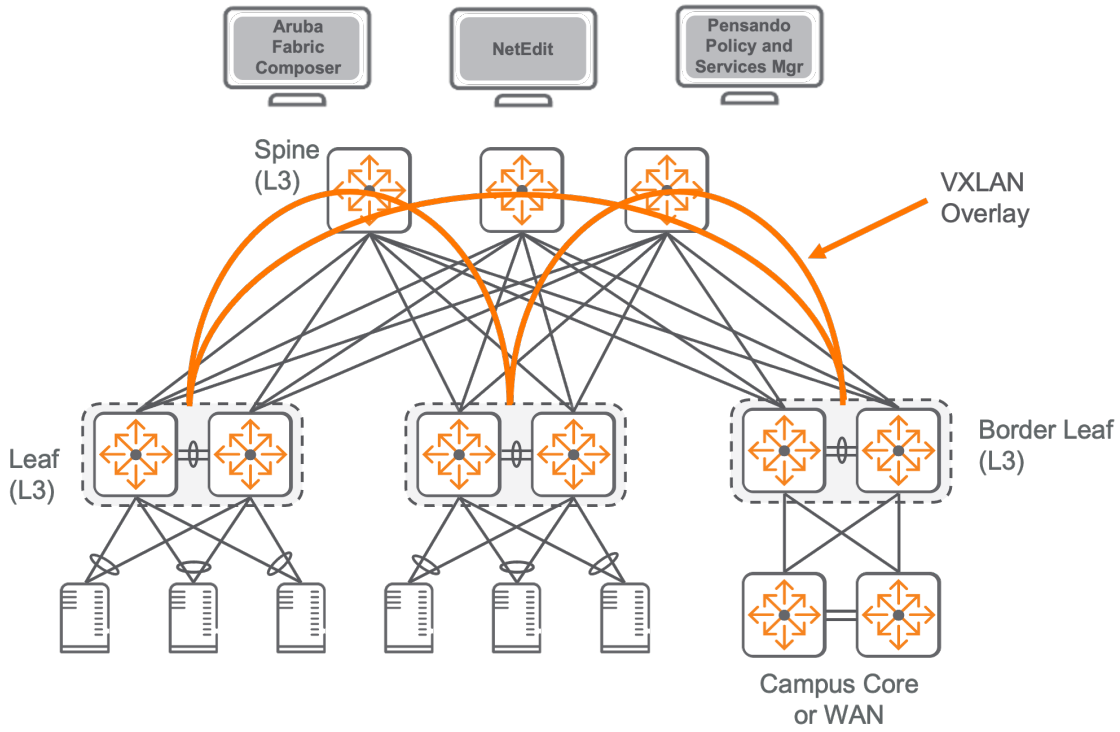

Figure 18: Spine and Leaf: Dual Top of Rack

Certain application environments do not require high availability at the individual computing host. In this case, a single ToR switch per rack provides a more cost-effective data center network. In this type of implementation, the number of computing hosts deployed per rack should be kept low because a ToR switch under maintenance impacts connectivity to all computing hosts in the rack. The following topology shows a single ToR design with two spine switches.

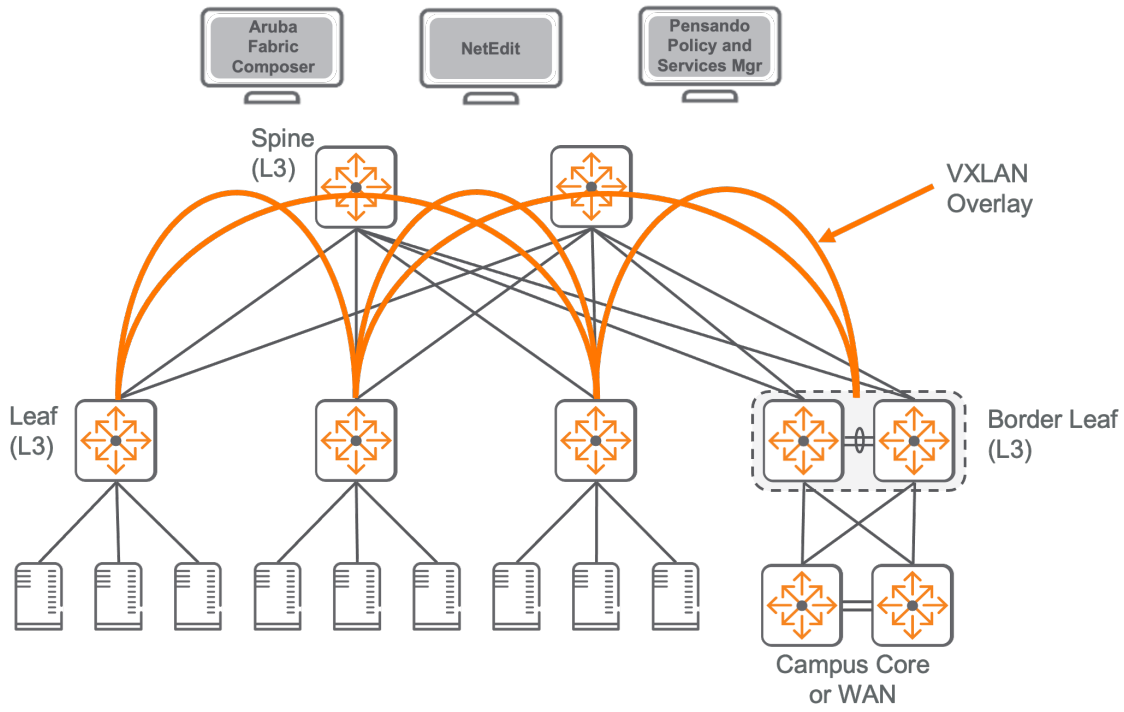

Figure 19: Spine and Leaf: Single Top of Rack

Reference Architecture Components Selection

The following section provides guidance for hardware selection based on your computing host, availability, and bandwidth requirements.

Aruba CX Data Center Switches

The Aruba CX portfolio offers three models of fixed configuration data center switches. The CX 10000 and 8325 models offer high port density, while the CX 8360 model offers a variety of port configurations for small and medium spine-and-leaf topologies. All models offer the following data center switching capabilities:

- High-speed, fully distributed architecture with line-rate forwarding
- High availability and in-service ToR upgrades with VSX
- Cloud-native and fully programmable modern operating system built on a microservices architecture
- Error-free network configuration with software-defined orchestration tools
- Distributed analytics and guided troubleshooting to provide full visibility and rapid issue resolution
- Hot-swappable and redundant load-sharing fans and power supplies
- Power-to-port and port-to-power cooling options for different data center designs
- Jumbo frame support for 9198 byte frames

- Advanced Layer 2 and Layer 3 features to support VXLAN spine and leaf with MP-BGP / EVPN control plane
- Distributed active gateways to support host mobility.

The Aruba CX 10000 distributed services switch supports additional features to consider when selecting a leaf switch model. The onboard Pensando DPU currently supports inline stateful firewall enforcement and enhanced traffic visibility.

Future functions will include encryption services, DDoS protection, load-balancing, and NAT.

Spine Switches

The Aruba ESP data center reference architecture is built around two 1RU high-density spine switches with QSFP ports capable of 40/100 GbE speeds.

- The Aruba CX 8325 can support up to 32 computing racks in a single ToR switch topology or up to 16 computing racks in a dual ToR switch topology.
- The Aruba CX 8360 can support up to 12 computing racks in a single ToR switch topology or up to six computing racks in a dual ToR switch topology.

The primary function of spine switches is making routing decisions for the overlay. The primary design considerations when choosing a spine switch are:

- Port density
- Ports speeds
- Routing table sizes.

Table 1: Spine Switches

SKU	Description	Maximum Rack Capacity
JL626A	8325: 32-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	32 single ToR / 16 dual ToR
JL627A	8325: 32-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	32 single ToR / 16 dual ToR
JL708A	8360: 12-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	12 single ToR / 6 dual ToR
JL709A	8360: 12-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	12 single ToR / 6 dual ToR

Leaf Switches

There are three leaf-switch models to choose from in the Aruba ESP Data Center reference architecture. All models are 1RU ToR switches that support high-density racks using 10GbE copper or SFP+ ports. SFP ports on the Aruba CX 8360 model also support 10GBASE-T transceivers.

For redundant ToR designs, the high- and medium-density SKUs provide the minimum of four uplink ports required for a two-spine-switch topology. For non-redundant ToR design, medium- and low-density SKUs provide the minimum of two uplink ports required for a two-spine-switch topology.

The Aruba CX 10000 distributed services switch (DSS) adds inline firewall features typically provided by VM hypervisors attached to leaf switches or dedicated firewall appliances attached to a services leaf. The Aruba CX 10000 switch should be selected when these features are required by downstream hosts or to meet other data center goals. DSS features are not available on other CX switch models. A mix of different ToR leaf switch models can connect to a common spine. The CX 10000, 8325, and 8360 can be installed in leaf racks that do not require distributed service switch features.

The following table summarizes the leaf SKUs available and their corresponding supported designs.

Table 2: Leaf Switches

SKU	Description	Rack Design	Spine Design
R8P13A	10000: 48-port 1/10/25 GbE SFP/SFP+/SFP28, 6-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	High-density / Dual ToR	2–4 switches
R8P14A	10000: 48-port 1/10/25 GbE SFP/SFP+/SFP28, 6-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	High-density / Dual ToR	2–4 switches
JL624A	8325: 48-port 1/10/25 GbE SFP/SFP+/SFP28, 8-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	High-density / Dual ToR	2–6 switches
JL625A	8325: 48-port 1/10/25 GbE SFP/SFP+/SFP28, 8-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	High-density / Dual ToR	2–6 switches
JL706A	8360: 48-port 100M / 1GbE / 10GbE 10GBASE-T, 4-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	High-density / Dual ToR	2 switches
JL707A	8360: 48-port 100M / 1GbE / 10GbE 10GBASE-T, 4-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	High-density / Dual ToR	2 switches
JL700A	8360: 32-port 1/10/25 GbE SFP/SFP+/SFP28, 4-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	Medium-density / Dual ToR	2 switches
JL701A	8360: 32-port 1/10/25 GbE SFP/SFP+/SFP28, 4-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	Medium-density / Dual ToR	2 switches
JL710A	8360: 24-port 1/10 GbE SFP/SFP+, 2-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	Medium-density / Single ToR	2 switches
JL711A	8360: 24-port 1/10 GbE SFP/SFP+, 2-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	Medium-density / Single ToR	2 switches
JL702A	8360: 16-port 1/10/25 GbE SFP/SFP+/SFP28, 2-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	Low-density / Single ToR	2 switches

SKU	Description	Rack Design	Spine Design
JL703A	8360: 16-port 1/10/25 GbE SFP/SFP+/SFP28, 2-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	Medium-density / Single ToR	2 switches

Out-of-Band Management Switches

The Aruba ESP data center reference architecture uses a management LAN built on dedicated switching infrastructure to ensure reliable connectivity to data center infrastructure for automation, orchestration, and traditional management access. The following table lists the recommended switch models.

Table 3: Out-of-Band Management Switches

SKU	Description	Host ports
JL667A	Aruba CX 6300F 48-port 1 GbE and 4-port SFP56 Switch	48
JL668A	Aruba CX 6300F 24-port 1 GbE and 4-port SFP56 Switch	24
JL663A	Aruba CX 6300M 48-port 1 GbE and 4-port SFP56 Switch	48
JL664A	Aruba CX 6300M 24-port 1 GbE and 4-port SFP56 Switch	24
JL724A	Aruba 6200F 24G 4SFP+ Switch	24
JL726A	Aruba 6200F 48G 4SFP+ Switch	48
JL678A	Aruba 6100 24G 4SFP+ Switch	24
JL676A	Aruba 6100 48G 4SFP+ Switch	48

Aruba Fabric Composer

Aruba Fabric Composer (AFC) is offered as a self-contained ISO or virtual machine OVA and can be installed in both virtual and physical host environments as a single instance or as a high-availability, three-node cluster. AFC is available as an annual per-switch software subscription.

AFC supports the Aruba CX 10000, 8325, and 8360 series switches recommended for spine and leaf roles. It also supports Aruba CX 6300, 6400, and 8400 series switches.

Ordering information for AFC is provided at the end of the solutions overview.

Pensando Policy and Services Manager

The Pensando Policy and Services Manager (PSM) runs as a virtual machine OVA on a host. PSM requires vCenter for installation. It is deployed as a high-availability, quorum-based cluster of three VMs.

PSM supports Aruba CX 10000 series switches. Management of PSM is integrated into AFC.

PSM can be downloaded from the Aruba Support Portal. Entitlement to PSM is included with the purchase of an Aruba CX 10000 switch by adding the R9H25AAE SKU.

NetEdit

NetEdit runs as a VM OVA on a host. Aruba NetEdit is available from the Aruba Service Portal. Customers must visit the Aruba Airheads Community and create an Airheads account in order to download the NetEdit software.

Ordering information for Aruba NetEdit is provided at the end of this data sheet.

Reference Architecture Physical Layer Planning

The following section provides guidance for planning the physical layer of data center switches.

Cables and Transceivers

Refer to the following documents to ensure that supported cables and transceivers are selected when planning physical connectivity inside the data center:

HPE Server Networking Transceiver and Cable Compatibility Matrix

ArubaOS-Switch and ArubaOS-CX Transceiver Guide

Port Speed Groups

For ToR configurations that require server connectivity at multiple speeds, it is important to note that setting the speed of a port might require adjacent ports to operate at that same speed.

Aruba CX 8325 and Aruba CX 10000 switches have a default speed of 25GbE. Changing the speed to 10GbE will impact groups of 12 ports on the Aruba CX 8325 and groups of four ports on the Aruba CX 10000. Aruba CX 8360 switches allow individual ports to operate at different speeds without impacting adjacent ports unless Media Access Control security (MACsec) is in use. Ports configured to use MACsec must all be configured to operate at the same speed.

Split Ports

Breakout cables can be used to split a 40 Gb/s or 100 Gb/s port into four lower-speed connections (4x10 Gb/s and 4x25 Gb/s). Refer to the ArubaOS-Switch and ArubaOS-CX Transceiver Guide when selecting supported breakout cables and switch support for split ports.

Media Access Control Security

MACsec is a standard defined in IEEE 802.1AE that extends standard Ethernet to provide frame-level encryption on point-to-point links. This feature is typically used in environments where additional layers of data confidentiality are required or where it is impossible to physically secure the network links between systems. Refer to the following table for details of MACsec support in the Aruba switching portfolio:

Table 4: MACsec Support in Aruba Switches

SKU	Description	Supported Ports
JL700A	8360: 32-port 1/10/25 GbE SFP/SFP+/SFP28, 4-port 40/100 GbE QSFP+/QSFP28, port-to-power airflow	1–4 SFP+/SFP28
JL701A	8360: 32-port 1/10/25 GbE SFP/SFP+/SFP28, 4-port 40/100 GbE QSFP+/QSFP28, power-to-port airflow	1–4 SFP+/SFP28

Reference Architecture Capacity Planning

The following section provides capacity planning guidance for the Aruba ESP data center spine-and-leaf reference architecture.

Bandwidth Calculations

A spine-and-leaf network design provides maximum flexibility and throughput for Aruba ESP data center implementation. To achieve the greatest level of performance, a spine-and-leaf topology can be designed for zero oversubscription of bandwidth. This results in a data center network that will never be congested because the bandwidth available to hosts is equal to the bandwidth between leaf-and-spine switches.

A significant advantage of a spine-and-leaf design is the ability to add capacity as needed simply by adding additional spine switches and/or increasing the speed of the uplinks between leaf-and-spine switches. A rack with 40 dual-homed servers with 10 GbE NICs could theoretically generate a total load of 800G of traffic. For that server density configuration, a 1:1 (non-oversubscribed) fabric could be built with four spine switches using 4x100 GbE links on each. In practice, most spine-and-leaf topologies are built with server-to-fabric oversubscription ratio between 2.4:1 and 6:1.

Network and Compute Scaling

The Aruba ESP data center reference architecture provides enough capacity for most deployments. With distributed gateways and symmetric IRB forwarding, the MAC and ARP tables are localized to directly attached computing nodes and are not impacted by the number of racks. The amount of IP prefixes is a function of the total number of nodes and fabric links, as well as the number of physical and/or virtualized servers. The border leaf is typically the node with the highest control plane load since it handles both internal and external connections. Route summarization is a good practice to reduce the redistribution of IP prefixes among domains.

The Aruba ESP data center reference architecture was thoroughly tested in an end-to-end solution environment that incorporates best-practice deployment recommendations, applications, and load profiles that represent production environments.

Refer to the product data sheets on Aruba Campus Core and Aggregation Switches for detailed specifications not included in this guide. The following table provides validated multi-dimensional profiles for spine-and-leaf design capacity planning.

Table 5: Validated Multi-Dimensional Profiles

Feature	8325 Leaf	8360 Leaf	8325 Spine	8360 Spine
Host scale—IPv4/ARP	30,000	50,000	N/A	N/A
Host scale—IPv6/ND	15,000	50,000	N/A	N/A
Routing—IPv4 routes	10,000	16,000	72,000	100,000
Routing—IPv6 routes	1000	8000	20,000	100,000
Routing—OSPF neighbors	4	4	128	64
VXLAN—overlay VRFs (Layer 3 VNI)	32	32	N/A	N/A
VXLAN—host VLANs (Layer 2 VNI)	1024	512	N/A	N/A
Active gateway SVIs	1000	512	N/A	N/A

Summary

Data center networks are changing rapidly. The most pressing challenge is maintaining operational stability, security, and visibility while placing computing and storage resources where they best serve users. In addition, data center teams are asked to support the rapid pace of DevOps environments, including connecting directly with public cloud infrastructure.

Given the rapidly changing landscape for data center requirements, it is critical that network and system engineers have the tools they need to simplify and automate complex infrastructure configurations.

The Aruba Networks ESP Data Center is built on technology that provides tools to transform the data center into a modern, agile services delivery platform that satisfies the requirements of distributed or centralized organizations of any size.

ArubaOS-CX simplifies operations and maintenance with a common switch OS across campus, branch, and data center, managed from the cloud or on-premises, and backed by AI that provides best-practice guidance throughout the network lifecycle.

What’s New in This Version

The following changes were made since Aruba last published this guide: - Expanded policy layer design guidance. - Expanded PSM and Aruba CX 10000 details.

www.arubanetworks.com

3333 Scott Blvd. Santa Clara, CA 95054
1.844.472.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550